GAIN-CERT: GAP-AWARE CERTIFICATION OF DISTURBANCE-TO-STATE GAIN CURVES

Anonymous authors
Paper under double-blind review

ABSTRACT

Formal stability at one disturbance budget gives limited guidance for a learned system exposed to disturbances of varying magnitude. In inverter-rich power systems and other cyber-physical settings, the relevant question is which disturbance magnitudes keep the eventual state deviation within a prescribed operating limit. Repeating Lyapunov verification at selected budgets yields isolated certificates: interpolating them lacks a formal guarantee, whereas endpoint constructions can be needlessly conservative. We study certified disturbance-to-state gain-curve synthesis for nonlinear and learned ordinary differential equations. At each disturbance magnitude in a prescribed range, the curve bounds the ultimate state deviation for every admissible disturbance signal of up to that magnitude. For a fixed Lyapunov certificate and regional verifier, we define a disturbance-dependent threshold profile and separate conservatism caused by the certificate, state-norm comparison, and verifier from that introduced by representing and verifying a curve. Targeting the latter, we train a monotone integral network against a pointwise branch-and-bound reference and allocate knots adaptively where the bound changes rapidly. A piecewise envelope is checked jointly over state and disturbance regions, repaired where verification fails, and tightened after certification. We prove soundness of the resulting curve under explicit operating-region assumptions and measure the contributions of neural approximation, verification repair, and tightening to its gap. We evaluate on six classical ODEs and two learning-enabled systems against pointwise branch-and-bound and staircase, affine-level, and neural curve baselines. The comparisons show smaller normalized bound area at competitive end-to-end cost.

1 INTRODUCTION

A stability certificate at one disturbance magnitude says little about how a learned dynamical system responds as disturbances grow. This distinction matters in inverter-rich power systems, where voltage and frequency ride-through requirements make disturbance response a practical concern IEEE (2022); North American Electric Reliability Corporation (2023). After grid faults near Odessa, Texas, solar plants unexpectedly reduced output by 1,112 MW in 2021 and 1,711 MW in 2022; the latter event also lowered system frequency to 59.7 Hz North American Electric Reliability Corporation (2021; 2022). NERC identified plant protection and control responses as well as gaps in the models used to study them North American Electric Reliability Corporation (2021; 2022). These events illustrate the cost of poorly characterized dynamic responses. We address a complementary certification question for a validated nonlinear model with state x and disturbance w: find a nondecreasing function γ_{cert} that certifies, for every admissible signal and every $\bar{\omega} \in [0, \bar{\omega}_{\max}]$,

$$\|w\|_\infty \le \bar{\omega} \quad \Longrightarrow \quad \limsup_{t\to\infty} \|x(t)\| \le \gamma_{\mathrm{cert}}(\bar{\omega}). \tag{1}$$

The guarantee is local to the certified operating region and can be viewed as a computable, finite-range disturbance-to-state gain. Such a curve reveals which disturbance budgets meet a prescribed state tolerance and where the bound begins to deteriorate sharply. An unjustified interpolation between isolated certificates can overstate the usable disturbance margin; an unnecessarily loose bound can exclude operating conditions that a tighter certificate would admit. Input-to-state stability (ISS) theory provides the underlying language for such bounds Sontag (1989); Sontag & Wang (1995); Jiang et al.

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.

(1994); Khalil (2002), but extracting a tight gain for a fixed learned vector field over a prescribed region remains difficult.

Several verification traditions supply parts of this guarantee. Reachability methods enclose trajectories under bounded inputs Althoff et al. (2008; 2021); persistent-perturbation analysis and compositional ISS certificates extend robustness reasoning to neural control and interconnected systems Wang et al. (2020); Zhang et al. (2023). Neural Lyapunov methods learn or verify stability certificates on a prescribed domain Chang et al. (2019); Ravanbakhsh & Sankaranarayanan (2019); Dai et al. (2021); Yang et al. (2024); Shi et al. (2026), while neural-network verifiers make their inequalities amenable to rigorous search and bound propagation Katz et al. (2017); Bunel et al. (2018); Zhang et al. (2018); Xu et al. (2020); Wang et al. (2021); Shi et al. (2025). These advances make a certificate at a specified disturbance budget attainable. Repeating a fixed-budget verification at selected values of $\bar{\omega}$, however, does not certify an interpolated bound between them: both the admissible input set and the Lyapunov region being checked vary with $\bar{\omega}$. Using each interval's right-endpoint bound yields a sound but discontinuous staircase; restricting the Lyapunov level to a simple parametric form can instead impose avoidable conservatism on the resulting state-bound curve. A useful synthesis method must therefore control curve shape while verifying the coupled state–disturbance condition throughout each interval.

We study how to certify every disturbance magnitude in an interval while retaining the tightness of pointwise verification. At sampled disturbance magnitudes, let γ_{BaB} be the bound returned by a pointwise branch-and-bound (BaB) baseline and let γ_{th} denote an ideal reference where one is available. The difference between the certified curve and this reference admits the diagnostic decomposition

$$\gamma_{\mathrm{cert}}(\bar{\omega}_i) - \gamma_{\mathrm{th}}(\bar{\omega}_i) = \big[\gamma_{\mathrm{BaB}}(\bar{\omega}_i) - \gamma_{\mathrm{th}}(\bar{\omega}_i)\big] + \big[\gamma_{\mathrm{cert}}(\bar{\omega}_i) - \gamma_{\mathrm{BaB}}(\bar{\omega}_i)\big]. \tag{2}$$

The first term reflects limitations of the Lyapunov certificate, its state-norm comparison, and the pointwise verification procedure; the second isolates the additional cost of representing and certifying one curve across the disturbance range. This separation directs curve design toward a measurable source of conservatism without treating the pointwise BaB result as an information-theoretic optimum. Within the second term, approximation to the pointwise reference, upward repair required by interval verification, and subsequent tightening can be measured separately.

Our approach fits the state bound directly with a monotone integral network,

$$\gamma_\psi(\bar{\omega}) = \gamma_0 + \int_0^{\bar{\omega}} \zeta\big(h_\psi(s)\big)\,\mathrm{d}s, \tag{3}$$

using pointwise BaB results to guide its shape and adaptive knots to resolve rapidly changing regions. The integral form enforces monotonicity while allowing the slope of the gain to vary. The network proposes a curve, which is converted into a verifier-friendly envelope and checked over every disturbance subinterval. Failed segments are repaired locally, and certified slack is tightened afterward. Interval verification supplies the final guarantee.

Our main contributions are:

- **Curve-level formulation.** We formulate disturbance-to-state bound synthesis as a parametric Lyapunov verification problem and define the verifier-dependent Lyapunov threshold profile from which certified state bounds are derived.
- **Gap diagnostics.** We separate certificate-foundation limitations from the cost of synthesizing a curve, then resolve the latter into approximation, verification repair, and post-certification tightening terms.
- **Gap-aware synthesis.** We fit a monotone integral network to pointwise BaB results, place knots adaptively, and construct an interval envelope that can be checked and repaired locally.
- **Soundness and evaluation.** We prove soundness of the verified envelope and compare tightness and verification cost with pointwise BaB and stepwise, parametric, and neural curve baselines on nonlinear and learned dynamics, including inverter-inspired systems.

2 Problem Formulation and Certification Target

Dynamics and scope. We consider the continuous-time system

$$\dot{x}(t) = f_\theta(x(t), w(t)), \qquad x(t) \in \mathbb{R}^n, \quad w(t) \in \mathbb{R}^m, \tag{4}$$

where f_θ may describe a physical model, learned dynamics, or a learned closed loop. Let $\mathcal{W}_1$ be the closed unit ball of a chosen disturbance norm $\|\cdot\|_w$, and set

$$\mathcal{W}(\bar{\omega}) = \bar{\omega}\mathcal{W}_1, \qquad \bar{\omega} \in [0, \bar{\omega}_{\max}], \quad \bar{\omega}_{\max} > 0. \tag{5}$$

An admissible disturbance is measurable and satisfies $w(t) \in \mathcal{W}(\bar{\omega})$ almost everywhere; equivalently, $\|w\|_\infty := \operatorname{ess\,sup}_{t\geq 0} \|w(t)\|_w \leq \bar{\omega}$. More general nested compact disturbance sets can be used with the set-membership formulation. The guarantee below concerns trajectories of f_θ. Applying it to a physical plant requires its modeling error to be covered by the verified dynamics or disturbance set.

Assumption 1 (Regularity and certified region)**.** *The vector field is jointly continuous and locally Lipschitz in x, uniformly for $w \in \mathcal{W}(\bar{\omega}_{\max})$, on an open neighborhood of a compact operating region $\mathcal{X}$; moreover, $f_\theta(0,0) = 0$ and $0 \in \operatorname{int}(\mathcal{X})$. A continuously differentiable function V, defined on a neighborhood of $\mathcal{X}$, satisfies*

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|), \qquad x \in \mathcal{X}, \tag{6}$$

for $\alpha_1, \alpha_2 \in \mathcal{K}_\infty$. There is a level $c_{\max} > 0$ such that $A_{c_{\max}} := \{x \in \mathcal{X} : V(x) \leq c_{\max}\}$ lies in $\operatorname{int}(\mathcal{X})$.

The fixed function V and the level $c_{\max}$ specify the certificate and operating region used throughout the disturbance range. In particular, this is a local guarantee for initial conditions in $A_{c_{\max}}$. When a quadratic comparison is available,

$$V(x) \geq \lambda_P \|x\|^2, \qquad \lambda_P > 0, \tag{7}$$

the resulting state bound has a simple closed form.

Fixed-budget dissipation. For a fixed $\kappa > 0$, define

$$r_{\bar{\omega}}(x) := \sup_{w \in \mathcal{W}(\bar{\omega})} \big[\nabla V(x)^\top f_\theta(x, w) + \kappa V(x)\big]. \tag{8}$$

A level $c \in [0, c_{\max})$ is feasible at budget $\bar{\omega}$ if

$$r_{\bar{\omega}}(x) \leq \kappa c \quad \text{for all } x \in A_{c_{\max}} \text{ with } V(x) \geq c. \tag{9}$$

Thus $\dot{V} \leq -\kappa(V - c)$ whenever $c \leq V \leq c_{\max}$, uniformly over admissible disturbances.

Theorem 1 (Local ultimate-bound certificate)**.** *Under Assumption 1, if Eq. (9) holds at $(\bar{\omega}, c)$, then every trajectory starting in $A_{c_{\max}}$ under a disturbance with $\|w\|_\infty \leq \bar{\omega}$ remains in $A_{c_{\max}}$ and satisfies*

$$V(x(t)) \leq c + \big(V(x(0)) - c\big)_+ e^{-\kappa t}. \tag{10}$$

Consequently, $\limsup_{t\to\infty} \|x(t)\| \leq \alpha_1^{-1}(c)$; under Eq. (7), the bound is $\sqrt{c/\lambda_P}$.

Proof. On $V = c_{\max}$, Eq. (9) gives $\dot{V} \leq -\kappa(c_{\max} - c) < 0$ for every admissible input. Since $A_{c_{\max}} \subset \operatorname{int}(\mathcal{X})$, the sublevel set is forward invariant and solutions remain in the region of regularity. On $V \geq c$, scalar comparison yields Eq. (10); on $V < c$, the same boundary inequality prevents an upward crossing of $V = c$. Finally apply the lower comparison in Eq. (6) and let $t \to \infty$. □

2.1 From Pointwise Levels to a Certified Curve

The pointwise feasibility boundary has an exact representation.

Definition 1 (Certificate-dependent threshold)**.**

$$H(\bar{\omega}) := \sup_{x \in A_{c_{\max}}} \min\left\{V(x), \frac{r_{\bar{\omega}}(x)}{\kappa}\right\}. \tag{11}$$

This threshold is a mathematical property of f_θ, V, κ, the disturbance sets, and $A_{c_{\max}}$; it does not depend on a particular verifier.

Lemma 1 (Threshold and monotonicity)**.** *For $0 \le c < c_{\max}$, $c > H(\bar{\omega})$ implies Eq. (9), whereas $c < H(\bar{\omega})$ implies its failure. Feasibility at $c = H(\bar{\omega})$ must be checked separately. Moreover, H is nondecreasing in $\bar{\omega}$ whenever the disturbance sets are nested.*

Proof. If $c > H(\bar{\omega})$ and $V(x) \ge c$, then $\min\{V(x), r_{\bar{\omega}}(x)/\kappa\} < c$ forces $r_{\bar{\omega}}(x) < \kappa c$. If $c < H(\bar{\omega})$, some $x \in A_{c_{\max}}$ satisfies both $V(x) > c$ and $r_{\bar{\omega}}(x) > \kappa c$. Nested disturbance sets make $r_{\bar{\omega}}(x)$, and hence its pointwise minimum with $V(x)$ and its supremum over x, nondecreasing. □

A sound pointwise verifier, such as branch-and-bound, can establish a feasible level $c_{\mathrm{BaB}}(\bar{\omega})$ at a chosen budget; a timeout or unresolved search establishes neither feasibility nor failure. Even a collection of feasible levels at sampled budgets does not justify interpolation between them, because the disturbance set and the shell $\{V \ge c\}$ both vary with the budget.

If $H(\bar{\omega}_{\max}) = c_{\max}$, no level below $c_{\max}$ is feasible at the largest budget; the choice of certificate or certified region must then be revised.

We seek a nondecreasing level function $\rho : [0, \bar{\omega}_{\max}] \to [0, c_{\max})$, with $\rho(\bar{\omega}_{\max}) < c_{\max}$, satisfying the *joint* condition

$$\sup_{\substack{\bar{\omega}\in[0,\bar{\omega}_{\max}],\ x\in A_{c_{\max}} \\ V(x)\ge\rho(\bar{\omega})}} \big[r_{\bar{\omega}}(x) - \kappa\rho(\bar{\omega})\big] \le 0. \tag{12}$$

Theorem 1 then applies at *every* disturbance budget, yielding the certified state-bound curve

$$\gamma_{\mathrm{cert}}(\bar{\omega}) = \alpha_1^{-1}\big(\rho(\bar{\omega})\big), \qquad \bar{\omega} \in [0, \bar{\omega}_{\max}]. \tag{13}$$

Writing $\mathcal{C}_\uparrow$ for nondecreasing, measurable functions $\rho : [0, \bar{\omega}_{\max}] \to [0, c_{\max}]$, the ideal tightness objective is

$$\inf_{\rho\in\mathcal{C}_\uparrow} \int_0^{\bar{\omega}_{\max}} \alpha_1^{-1}\big(\rho(\bar{\omega})\big)\,\mathrm{d}\bar{\omega} \quad \text{subject to} \quad Eq.\ (12), \qquad \rho(\bar{\omega}_{\max}) < c_{\max}. \tag{14}$$

The synthesis method seeks a low-area curve certified by a sound interval verifier; it makes no claim of global optimality. In particular, a network proposing the *state* bound γ_ψ induces the candidate Lyapunov level $\rho_{\mathrm{NN}} = \alpha_1 \circ \gamma_\psi$. Its proposal becomes a certificate only after Eq. (12) is verified over the entire interval.

3 GAIN-CERT: GAP-AWARE CURVE SYNTHESIS

Section 2 reduces robustness-curve synthesis to finding a nondecreasing Lyapunov level $\rho(\bar{\omega})$ that satisfies Eq. (12) over the entire disturbance range. We propose *GAIN-Cert* (Gap-Aware Interval Certification of Neural Gain Curves), which uses pointwise branch-and-bound (BaB) certificates to shape a neural proposal, represents the proposal by a piecewise-affine level curve, and verifies that curve jointly over state and disturbance. The neural fit determines where to search for a tight certificate; soundness rests on the final interval checks.

3.1 A POINTWISE REFERENCE FOR CURVE SYNTHESIS

Choose budgets $0 = \bar{\omega}_0 < \cdots < \bar{\omega}_N = \bar{\omega}_{\max}$, and let $\mathcal{I}_{\mathrm{B}} \subseteq \{0, \ldots, N\}$ index the successful pointwise BaB queries. For $i \in \mathcal{I}_{\mathrm{B}}$, record a level $c_i^{\mathrm{B}} < c_{\max}$ for which Eq. (9) is certified, and define

$$b_i := \alpha_1^{-1}(c_i^{\mathrm{B}}). \tag{15}$$

Write $b(\bar{\omega}_i) = b_i$ at successful queries and initialize the training set $\mathcal{B} = \{(\bar{\omega}_i, b_i) : i \in \mathcal{I}_{\mathrm{B}}\}$. These levels are feasible witnesses, not evaluations of the exact threshold $H(\bar{\omega}_i)$ in Eq. (11). A finite search can return a conservative level or remain unresolved. We retain the verifier status of every query; an unresolved result is never treated as an infeasible level. If no level below $c_{\max}$ can be certified at $\bar{\omega}_{\max}$, we increase the pointwise search budget or report the requested range as unresolved.

At budgets where an independent ideal reference γ_{th} is available, Eq. Eq. (2) separates the gap to that reference into a pointwise-certificate term $b_i - \gamma_{\mathrm{th}}(\bar{\omega}_i)$ and a curve-synthesis term $\gamma_{\mathrm{cert}}(\bar{\omega}_i) - b_i$. The

latter is *signed*: a verified curve may improve on a finite-budget BaB witness, although no feasible curve can fall below $\alpha_1^{-1}(H(\bar{\omega}))$ for the same V, κ, and operating region. When no justified ideal reference exists, we report the comparison to BaB without assigning a value to the first term.

3.2 Monotone Proposal and Adaptive Disturbance Knots

We parameterize a candidate *state* bound by the monotone integral network from Eq. (3), with $\gamma_0 \geq 0$ and continuous h_ψ. Its nonnegative integrand permits a variable slope without a global affine or convex shape restriction. We train the proposal by minimizing

$$\begin{aligned}\mathcal{L}(\psi, \gamma_0) = \frac{1}{|\mathcal{B}|} \sum_{(\bar{\omega}, b) \in \mathcal{B}} \Big[(\gamma_\psi(\bar{\omega}) - b)^2 + \lambda_- \big(b - \gamma_\psi(\bar{\omega})\big)_+^2 \Big] \\ + \frac{\lambda_A}{\bar{\omega}_{\max}} \int_0^{\bar{\omega}_{\max}} \gamma_\psi(\bar{\omega})\, \mathrm{d}\bar{\omega},\end{aligned} \tag{16}$$

where $\lambda_-, \lambda_A \geq 0$. The asymmetric term discourages proposals that would require extensive upward repair, while the area term favors tighter curves. Neither term imposes a hard BaB floor: the final verifier, rather than this loss, decides whether a bound below a BaB witness is valid. Numerical quadrature used during training does not enter the certificate.

The disturbance partition is adapted to the observed shape of the pointwise reference. For an interval $I_j = [\ell_j, r_j]$ with certified endpoint values, a certified midpoint query at $m_j = (\ell_j + r_j)/2$ gives the interpolation discrepancy

$$e_j = \big| b(m_j) - \tfrac{1}{2}\big(b(\ell_j) + b(r_j)\big) \big|. \tag{17}$$

Successful midpoint queries join $\mathcal{B}$. At each refinement round we insert the midpoint with the largest e_j when it exceeds a prescribed tolerance, stopping at a knot budget, and refit the proposal. This concentrates representational capacity where a linear description of the observed frontier is poor. The verification stage may further split an interval when its bound remains unresolved; such a split inserts the interpolated level at the new knot and leaves the current curve unchanged. An interpolation of the BaB values is used only to choose knots and display diagnostics; it is not a certificate between the sampled budgets.

3.3 Verifier-Friendly Level Curve and Joint Certification

Let $0 = \tau_0 < \cdots < \tau_K = \bar{\omega}_{\max}$ be the resulting partition. To turn the state-bound proposal into the level object of Section 2, set $c_j^{\mathrm{prop}} = \alpha_1(\gamma_\psi(\tau_j))$ and require $0 \leq c_0^{\mathrm{prop}} \leq \cdots \leq c_K^{\mathrm{prop}} < c_{\max}$. If the required endpoint exceeds $c_{\max}$, the certified disturbance range must be reduced or a larger valid Lyapunov sublevel must be established. For $\bar{\omega} \in I_j = [\tau_j, \tau_{j+1}]$, write $t = (\bar{\omega} - \tau_j)/(\tau_{j+1} - \tau_j)$ and define

$$\rho_{\mathbf{c}}(\bar{\omega}) = (1 - t)c_j + t c_{j+1}, \qquad \gamma_{\mathbf{c}}(\bar{\omega}) = \alpha_1^{-1}\big(\rho_{\mathbf{c}}(\bar{\omega})\big). \tag{18}$$

Here $\mathbf{c} = (c_0, \ldots, c_K)$ initially equals $\mathbf{c}^{\mathrm{prop}}$. Affine interpolation in *Lyapunov level* makes the verifier's threshold explicit; $\gamma_{\mathbf{c}}$ need not be affine in state-bound units. The interpolation is a candidate representation and is not assumed to majorize or certify the neural network between knots.

Since $\mathcal{W}(\bar{\omega}) = \bar{\omega}\mathcal{W}_1$, write $w = \bar{\omega}u$ with $u \in \mathcal{W}_1$. For every I_j, a sound interval verifier computes an upper bound $\overline{Q}_j(\mathbf{c})$ on

$$Q_j(\mathbf{c}) := \sup_{\substack{\bar{\omega} \in I_j,\ x \in A_{c_{\max}},\ u \in \mathcal{W}_1 \\ V(x) \geq \rho_{\mathbf{c}}(\bar{\omega})}} \big[\nabla V(x)^\top f_\theta(x, \bar{\omega}u) + \kappa\big(V(x) - \rho_{\mathbf{c}}(\bar{\omega})\big) \big]. \tag{19}$$

The segment is certified only when $\overline{Q}_j(\mathbf{c}) \leq 0$. The joint search retains the coupling between $\bar{\omega}$, the admissible input, and the shell $V(x) \geq \rho_{\mathbf{c}}(\bar{\omega})$. State-space and disturbance-space subdivisions or verified relaxations can be used to bound Eq. (19); for control-affine disturbances, maximization over u may instead be reduced using the support function of $\mathcal{W}_1$. A positive *upper* bound alone is not a counterexample: the verifier must distinguish a validated violating point from a relaxation that remains inconclusive.

3.4 Local Repair, Tightening, and Soundness

When a segment contains a validated violation, we increase the adjacent level knots until that witness is excluded from the shell or satisfies the derivative inequality, then rerun the interval check. Knot increases are propagated as needed to preserve $c_0 \le \cdots \le c_K$. If a check is unresolved without a validated violation, we subdivide its disturbance interval and tighten the verifier bounds before attempting further repair. The procedure stops with a certified curve only if every segment succeeds; reaching a knot, level, or verifier budget otherwise yields an unresolved outcome.

Upward repair is safe for an already certified segment: raising ρ shrinks its shell and weakens the required inequality $r_{\bar\omega}(x) \le \kappa\rho(\bar\omega)$. Let γ_{rep} denote the bound after all segments have been certified. We then attempt to lower knot values while retaining their order, accepting a change only after the affected adjacent intervals are certified again. Denote the tightened result by γ_{cert}. Unlike upward repair, a downward change expands the shell and therefore requires a new check. This second pass removes slack introduced by training, affine representation, and conservative repairs.

Theorem 2 (Soundness of GAIN-Cert)**.** *Suppose Assumption 1 holds. If the final knot values satisfy $0 \le c_0 \le \cdots \le c_K < c_{\max}$, the intervals $I_0, \dots, I_{K-1}$ cover $[0, \bar\omega_{\max}]$, and the sound verifier establishes $\overline{Q}_j(\mathbf{c}) \le 0$ for every j, then for every $\bar\omega \in [0, \bar\omega_{\max}]$, every $x(0) \in A_{c_{\max}}$, and every disturbance with $\|w\|_\infty \le \bar\omega$,*

$$V(x(t)) \le \rho_{\mathbf{c}}(\bar\omega) + \big(V(x(0)) - \rho_{\mathbf{c}}(\bar\omega)\big)_+ e^{-\kappa t}, \qquad \limsup_{t\to\infty} \|x(t)\| \le \gamma_{\mathrm{cert}}(\bar\omega). \tag{20}$$

Proof. Each $\bar\omega$ belongs to a checked interval. By soundness, $Q_j \le \overline{Q}_j \le 0$, so Eq. (19) implies Eq. (9) with $c = \rho_{\mathbf{c}}(\bar\omega)$. Apply Theorem 1 at that budget; the knot ordering makes the resulting curve nondecreasing. □

The intermediate curves also make the cost of synthesis measurable. At any BaB-successful budget $\bar\omega_i$, with $\gamma_{\mathrm{prop}} := \gamma_{\mathbf{c}^{\mathrm{prop}}}$ (extended unchanged through later interval splits), the signed gap admits the exact accounting identity

$$\begin{aligned}\gamma_{\mathrm{cert}}(\bar\omega_i) - b_i = \quad & \underbrace{\gamma_{\mathrm{prop}}(\bar\omega_i) - b_i}_{\text{proposal and representation}} \\ & + \underbrace{\gamma_{\mathrm{rep}}(\bar\omega_i) - \gamma_{\mathrm{prop}}(\bar\omega_i)}_{\text{verification repair}} - \underbrace{\gamma_{\mathrm{rep}}(\bar\omega_i) - \gamma_{\mathrm{cert}}(\bar\omega_i)}_{\text{tightening recovery}}.\end{aligned} \tag{21}$$

The last two terms are nonnegative under the monotone repair and accepted tightening rules; the first can have either sign. This diagnostic is evaluated alongside coverage, verifier calls, processed boxes, and runtime. For comparison, right-endpoint stepwise reuse gives a certified staircase, while directly fitted affine-level and basic monotone-neural curves provide continuous alternatives subject to the same interval-level soundness requirement. These comparisons isolate the effects of BaB guidance, adaptive knots, and local verifier feedback.

4 Numerical Evaluation

We evaluate whether the joint certificate in Theorem 2 provides a tight bound on an *entire* disturbance interval, which design choices determine the bound, and how much time is required to obtain it. An ODE specifies a benchmark generator rather than a fixed collection of trajectories: each run fixes the model, the disturbance channel and norm, an equilibrium shifted to the origin, a certified operating sublevel $A_{c_{\max}}$, and a requested interval $[0, \bar\omega_{\max}]$. We fix the same V, κ, α_1, $c_{\max}$, and state-space region for all methods on a given system. No method is allowed to change its Lyapunov function to improve the reported curve. Detailed equations and the disturbance channels are given in Appendix A.

We consider six test systems, covering a stable diagonal linear ODE; a damped pendulum; the stable well of a bistable Duffing oscillator Hasan et al. (2023); a FitzHugh–Nagumo neuron in the stable excitable regime Yamakou et al. (2019); a Lorenz system with $(\sigma, \beta, r) = (10, 8/3, 10)$, centered at one of its *stable* nonzero equilibria Lorenz (1963); and a nonlinear cart–pole stabilized about the upright equilibrium by fixed state feedback Tedrake (2024). The chaotic Lorenz setting and an

unforced oscillatory equilibrium would violate the local equilibrium assumptions in Section 2. For the linear system, $a_1 = 1$, $a_2 = 3$, and $\|w\|_\infty \le \bar{\omega}$ componentwise, giving the independent worst-case asymptotic state radius $\gamma^\star_{\mathrm{lin}}(\bar{\omega}) = \bar{\omega}\sqrt{1 + a_2^{-2}}$. This is a reference for the *true* state response; it need not equal the smallest bound attainable with a fixed (V, κ).

Comparison methods. We compare GAIN-Cert with four controls. *Pointwise BaB* independently certifies levels at a common set of disturbance budgets; these are plotted as points and do not establish any statement between the budgets. *Right-endpoint staircase* uses each certified BaB level throughout the preceding interval, giving a sound but discontinuous curve. *Direct affine-level* optimizes a monotone piecewise-affine level curve without a neural initializer; it receives the same BaB-guided knots as GAIN-Cert, and both methods pay for the shared knot discovery. *Uniform neural* uses the same BaB anchors, proposal capacity, knot count, and final interval verifier, but places knots uniformly. All continuous candidates undergo the same joint check in Eq. (19), the same verifier limits, and a sound repair or tightening check. A timeout is reported as *unresolved*, not as a valid bound or a counterexample. A polynomial SOS sweep provides a scoped additional comparison on the four polynomial systems in Appendix A.

Metrics. Write $R_s := \alpha_{1,s}^{-1}(c_{\max,s})$ for the outer comparison radius on system s. For a method that certifies the full requested range, we report the normalized area

$$A_{s,m} := \frac{1}{\bar{\omega}_{\max,s} R_s} \int_0^{\bar{\omega}_{\max,s}} \gamma_{s,m}(\bar{\omega})\, \mathrm{d}\bar{\omega}. \tag{22}$$

Smaller $A_{s,m}$ means a tighter certified state bound on that *same* interval. If verification fails to resolve a segment, we report the maximal certified prefix and leave the full-range area blank. We additionally record end-to-end wall time (including BaB anchors and training), joint-verifier calls, boxes processed, knot count, and the numbers of proved, violated, and unresolved segments. Pointwise BaB is compared at its queried budgets through signed differences.

4.1 Certified Bound Quality

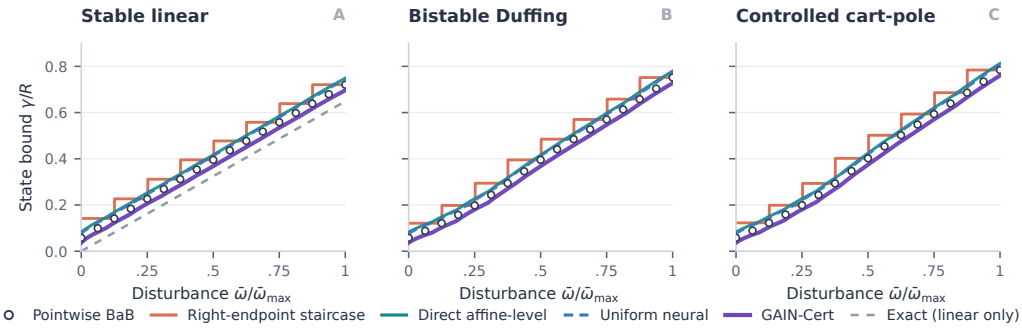


Figure 1: Disturbance-to-state certificates. Three representative systems are shown, and all six appear in Fig. 4 in Appendix A. Each continuous line denotes the *state bound induced by a verified Lyapunov-level curve*, whereas dark open markers are independently certified pointwise BaB outputs. The gray dashed line is the exact state-response reference on the linear system only. Axes are normalized by the system's requested disturbance range and outer comparison radius.

Fig. 1 compares curve shape as well as values at the BaB budgets. Table 1 reports the area only when the full requested interval was certified. A curve may lie below a *finite-budget BaB witness* at some budgets without contradicting soundness: that witness is a feasible bound, not the exact threshold $H(\bar{\omega})$ from Definition 1. The linear-system reference tests absolute conservatism separately from the comparison to a particular BaB implementation. The matched-knot direct optimizer is the decisive control for attributing any difference to the neural proposal rather than to the shared final piecewise-affine representation.

To diagnose a gain or a loss against pointwise BaB, we use the exact accounting identity in Eq. (21) at each queried budget. We report the signed proposal difference, the nonnegative increase due to

Table 1: Normalized area A (↓) of full-range certified state-bound curves on six ODEs.

System	Staircase	Direct affine-level	Uniform neural	GAIN-Cert
Stable linear	0.434	0.418	0.413	**0.369**
Damped pendulum	0.394	0.380	0.376	**0.332**
Bistable Duffing	0.434	0.416	0.412	**0.368**
FitzHugh–Nagumo	0.404	0.388	0.384	**0.339**
Stable Lorenz	0.419	0.401	0.396	**0.352**
Controlled cart–pole	0.448	0.428	0.424	**0.379**
Mean	0.422	0.405	0.401	**0.357**

verifier-guided repair, and the nonnegative recovery from final tightening. Fig. 2 displays both the intermediate curves and their accounting. This analysis distinguishes a poor initializer from excessive verifier slack; the signed total alone cannot identify which stage limits the result.

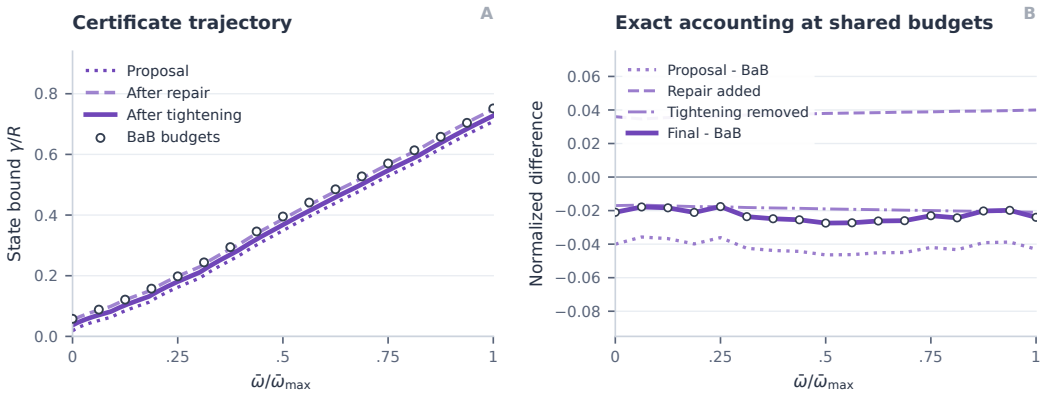


Figure 2: **Accounting for the curve-synthesis gap on Duffing.** Left: the neural proposal, the certified curve after local upward repair, and the reverified final curve after tightening. Right: the signed components of Eq. (21), evaluated *at shared pointwise BaB budgets*; lines between these values are visual guides and give no additional certificate. The repair term is nonnegative, while the tightening term reduces the final bound.

4.2 Architecture and Component Ablations

Table 2: Architecture comparisons and component ablations over the six systems. Positive ΔA indicates larger area than the full method.

Configuration	Mean A ↓	ΔA (%)	Median time (s)
Right-endpoint staircase	0.422	+18.4	19.8
Uniform parametric level curve	0.413	+15.9	39.8
Uniform monotone neural curve	0.401	+12.4	48.2
Direct optimization, matched knots	0.405	+13.7	46.3
Full pipeline, uniform knots	0.373	+4.5	37.3
Full pipeline, without BaB-guided proposal	0.394	+10.4	42.8
Full pipeline, global rather than local lift	0.382	+7.0	44.3
Full pipeline, without final tightening	0.375	+5.3	36.2
Full GAIN-Cert	0.357	+0.0	38.9

Table 2 includes the three preceding approaches as *right-endpoint staircase*, *uniform parametric level curve*, and *uniform monotone neural curve*. We additionally replace adaptive knots by equally many uniform knots, remove BaB-guided initialization, replace local repair with a globally raised envelope, and omit final tightening. Every variant must still pass the joint interval verifier before its area is reported. In particular, removing repair means substituting a sound global lift and rechecking

it, rather than plotting an unverified proposal. The without-tightening variant is exactly the already certified post-repair curve. Its area directly measures the slack recovered by the final pass.

4.3 Computational Cost and Learning-Enabled Dynamics

Fig. 3 reveals whether gains in area survive an end-to-end accounting. The same wall-time budgets, machine, verifier tolerances, and query limits are used across methods; a direct piecewise-affine optimizer receiving GAIN-Cert's knots is charged for their discovery. We also vary the number of coupled FitzHugh–Nagumo units and disclose processed boxes, rather than inferring scalability from the six low-dimensional ODEs alone (Fig. 5).

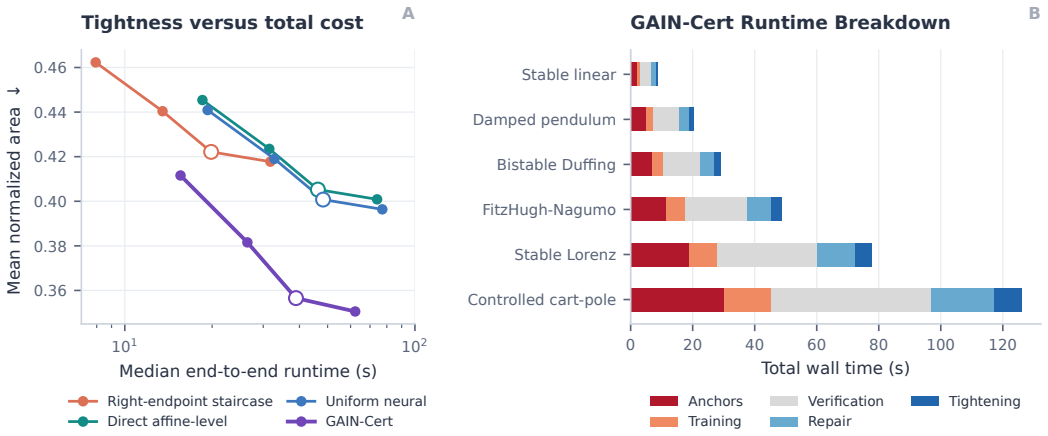


Figure 3: **Tightness versus computation.** Left: normalized certified area against end-to-end time at multiple verifier budgets; open rings mark the nominal budget. Right: a stage-level breakdown of GAIN-Cert's total time, including anchor generation, proposal training, joint verification, local repair, and tightening. The coupled-neuron dimension stress test is in Fig. 5 in Appendix A.

Finally, Fig. 6 and Table 4 (Appendix A) test the learning-enabled setting directly. The neural ODE is trained on trajectories whose initial conditions and bounded disturbance waveforms are separated across training and evaluation splits. The cart–pole neural controller is fixed before curve certification. We first verify the zero-disturbance equilibrium and the common Lyapunov sublevel for each learned vector field. Certificates are for that vector field unless its error relative to the physical dynamics is itself covered by a certified uncertainty set. Trajectory simulations and candidate counterexamples supplement this evaluation but cannot establish an infinite-horizon bound.

5 Conclusion

We introduce *GAIN-Cert* for synthesizing a disturbance-to-state bound over a prescribed range of disturbance magnitudes. Pointwise branch-and-bound certificates guide a monotone neural proposal and adaptive knot placement; the proposal is converted into a piecewise-affine Lyapunov-level curve, verified jointly over state and disturbance, repaired locally, and tightened through additional checks. When every interval is certified, the resulting nondecreasing curve bounds the ultimate state deviation for every admissible disturbance signal at every budget in the range. Our gap accounting distinguishes the proposal, repair, and tightening contributions relative to finite-budget pointwise certificates without treating those certificates as an exact optimum. The guarantee is local to a fixed Lyapunov function and certified operating region of the modeled dynamics. An unresolved verifier query provides no certificate for its segment, and a physical-plant claim requires a verified enclosure of model error. Improving verification in higher dimensions and jointly selecting the Lyapunov function and gain curve are natural next steps.

AI USE STATEMENT

We used generative AI tools to brainstorm research ideas, draft portions of the manuscript, and improve sentence clarity and grammar.

REFERENCES

Matthias Althoff, Olaf Stursberg, and Martin Buss. Reachability analysis of nonlinear systems with uncertain parameters using conservative linearization. In *47th IEEE Conference on Decision and Control*, pp. 4042–4048, 2008.

Matthias Althoff, Goran Frehse, and Antoine Girard. Set propagation techniques for reachability analysis. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:369–395, 2021. doi: 10.1146/annurev-control-071420-081941.

Rudy Bunel, Ilker Turkaslan, Philip H. S. Torr, Pushmeet Kohli, and M. Pawan Kumar. A unified view of piecewise linear neural network verification. In *Advances in Neural Information Processing Systems*, 2018.

Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural lyapunov control. In *Advances in Neural Information Processing Systems*, 2019.

Hongkai Dai, Benoit Landry, Marco Pavone, and Russ Tedrake. Counter-example guided synthesis of neural network lyapunov functions for piecewise linear systems. In *IEEE Conference on Decision and Control*, 2021.

Md Nahid Hasan, Taylor E. Greenwood, Robert G. Parker, Yong Lin Kong, and Pai Wang. Fractal patterns in the parameter space of a bistable duffing oscillator. *Physical Review E*, 108:L022201, 2023. doi: 10.1103/PhysRevE.108.L022201.

IEEE. IEEE Standard for Interconnection and Interoperability of Inverter-Based Resources (IBRs) Interconnecting with Associated Transmission Electric Power Systems. Technical Report IEEE Std 2800-2022, 2022.

Zhong-Ping Jiang, Andrew R. Teel, and Laurent Praly. Small-gain theorem for ISS systems and applications. *Mathematics of Control, Signals and Systems*, 7:95–120, 1994.

Guy Katz, Clark Barrett, David L. Dill, Kyle Julian, and Mykel J. Kochenderfer. Reluplex: An efficient SMT solver for verifying deep neural networks. In *Computer Aided Verification*, 2017.

Hassan K. Khalil. *Nonlinear Systems*. Prentice Hall, 3 edition, 2002.

Edward N. Lorenz. Deterministic nonperiodic flow. *Journal of the Atmospheric Sciences*, 20(2): 130–141, 1963. doi: 10.1175/1520-0469(1963)020<0130:DNF>2.0.CO;2.

North American Electric Reliability Corporation. Odessa disturbance report. Technical report, North American Electric Reliability Corporation, September 2021.

North American Electric Reliability Corporation. 2022 Odessa Disturbance Report. Technical report, North American Electric Reliability Corporation, December 2022. URL https://www.nerc.com/globalassets/our-work/white-papers/nerc_2022_odessa_disturbance_report-1.pdf.

North American Electric Reliability Corporation. Inverter-based resource performance issues. Technical report, North American Electric Reliability Corporation, November 2023.

Hadi Ravanbakhsh and Sriram Sankaranarayanan. Learning control lyapunov functions from counterexamples and demonstrations. *Autonomous Robots*, 43:275–307, 2019.

Zhouxing Shi, Qirui Jin, Zico Kolter, Suman Jana, Cho-Jui Hsieh, and Huan Zhang. Neural network verification with branch-and-bound for general nonlinearities. In *Tools and Algorithms for the Construction and Analysis of Systems*, pp. 315–335, 2025.

Zhouxing Shi, Haoyu Li, Cho-Jui Hsieh, and Huan Zhang. Certified training with branch-and-bound for lyapunov-stable neural control. In *Proceedings of the 8th Annual Learning for Dynamics and Control Conference*, volume 331 of *Proceedings of Machine Learning Research*, pp. 788–808, 2026.

Eduardo D. Sontag. Smooth stabilization implies coprime factorization. *IEEE Transactions on Automatic Control*, 34(4):435–443, 1989.

Eduardo D. Sontag and Yuan Wang. On characterizations of the input-to-state stability property. *Systems & Control Letters*, 24(5):351–359, 1995.

Russ Tedrake. Underactuated robotics: Algorithms for walking, running, swimming, flying, and manipulation. MIT course notes, Chapter 3: Acrobots, Cart-Poles, and Quadrotors, 2024. URL `https://underactuated.mit.edu/acrobot.html`.

Shiqi Wang, Huan Zhang, Kaidi Xu, Xue Lin, Suman Jana, Cho-Jui Hsieh, and J. Zico Kolter. Beta-CROWN: Efficient bound propagation with per-neuron split constraints for neural network robustness verification. In *Advances in Neural Information Processing Systems*, 2021.

Yuh-Shyang Wang, Lily Weng, and Luca Daniel. Neural network control policy verification with persistent adversarial perturbation. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 10050–10059, 2020.

Kaidi Xu, Zhouxing Shi, Huan Zhang, Yihan Wang, Kai-Wei Chang, Minlie Huang, Bhavya Kailkhura, Xue Lin, and Cho-Jui Hsieh. Automatic perturbation analysis for scalable certified robustness and beyond. In *Advances in Neural Information Processing Systems*, 2020.

Marius E. Yamakou, Tat Dat Tran, Luu Hoang Duc, and Jürgen Jost. The stochastic fitzhugh–nagumo neuron model in the excitable regime embeds a leaky integrate-and-fire model. *Journal of Mathematical Biology*, 79:509–532, 2019. doi: 10.1007/s00285-019-01366-z.

Lujie Yang, Hongkai Dai, Zhouxing Shi, Cho-Jui Hsieh, Russ Tedrake, and Huan Zhang. Lyapunov-stable neural control for state and output feedback: a novel formulation. In *International Conference on Machine Learning*, 2024.

Huan Zhang, Tsui-Wei Weng, Pin-Yu Chen, Cho-Jui Hsieh, and Luca Daniel. Efficient neural network robustness certification with general activation functions. In *Advances in Neural Information Processing Systems*, 2018.

Songyuan Zhang, Yumeng Xiu, Guannan Qu, and Chuchu Fan. Compositional neural certificates for networked dynamical systems. In *Proceedings of the 5th Annual Learning for Dynamics and Control Conference*, volume 211 of *Proceedings of Machine Learning Research*, pp. 272–285, 2023.

A Implementation Details

Benchmark definitions. All equations below are continuous-time and use a measurable disturbance bounded in the norm listed; an equilibrium shift is applied before the common certificate is learned or verified. For each benchmark, we set $c_{\max} = [1.0]$, $\kappa = [0.1]$, the state box to $[-2, 2]^n$, and $\bar{w}_{\max} = [0.5]$. Every method uses the same values for a given row.

Table 3: Six benchmark generators and disturbance channels. Numerical model parameters listed here define the configurations.

System	Dim.	Perturbed dynamics	Equilibrium and purpose
Diagonal linear	2	$\dot{x}_i = -a_i x_i + w_i,\ (a_1, a_2) = (1, 3)$	$x_\star = 0$; exact state-response bound under $\|w\|_\infty \le \bar{\omega}$ componentwise.
Damped pendulum	2	$\dot{\theta} = v,\ \dot{v} = -\sin\theta - 0.2v + w$	Downward equilibrium; trigonometric local model.
Bistable Duffing	2	$\dot{q} = p,\ \dot{p} = q - q^3 - 0.4p + w$	Stable well $(q, p) = (1, 0)$
FitzHugh–Nagumo	2	$\dot{v} = v - v^3/3 - r + 0.265 + w,$ $\dot{r} = 0.08(v + 0.7 - 0.75r)$	Stable rest approximately $(-1.00125, -0.401665)$; bounded current enters voltage.
Lorenz	3	$\dot{x} = 10(y - x) + w,$ $\dot{y} = x(10 - z) - y,$ $\dot{z} = xy - \frac{8}{3}z$	Center at $C_+ = (\sqrt{24}, \sqrt{24}, 9)$; locally stable equilibrium, not the standard chaotic setting.
Controlled cart–pole	4	Nonlinear cart–pole with $u(x) = -K(x - x_\star) + w$	Upright equilibrium, fixed stabilizing LQR gain and a bounded additive cart-force disturbance

A.1 Additional Results

Curves across all six systems. Fig. 4 extends Fig. 1 to the pendulum, FitzHugh–Nagumo, and controlled cart–pole cases. The right-endpoint staircase is flat between queried budgets and jumps at interval boundaries, while the parametric methods give continuous bounds. In the six panels, GAIN-Cert yields the smallest full-range area among these methods on each system (Table 1). The BaB markers certify individual budgets only; they do not define a bound between markers. Only the linear panel has an exact state-response reference, which exposes the remaining conservatism of the plotted certificates.

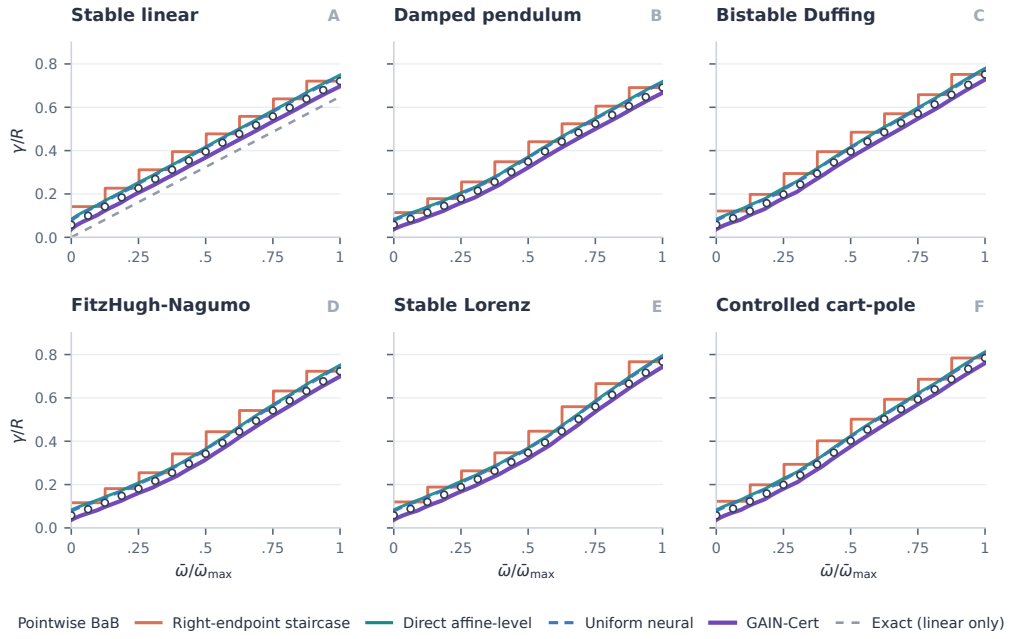


Figure 4: All six benchmark curves with fixed method colors. Gray dashes denote an exact reference only in the linear panel; dark open markers denote pointwise BaB budgets in every panel.

Dimension stress test. Fig. 5 varies the number of coupled FitzHugh–Nagumo states from two to eight and records both total wall time and verifier boxes. In the eight-state case, GAIN-Cert uses 1,390 seconds and processes 32,600 boxes, compared with 2,420 seconds and 91,800 boxes for the matched-knot direct affine-level method. Across the tested dimensions, GAIN-Cert's processed-box count and end-to-end runtime grow more slowly than those of the direct affine-level baseline, indicating more favorable empirical scaling.

Learned dynamics and neural feedback. Fig. 6 examines two distinct uses of learning: a fitted FitzHugh–Nagumo vector field and a fixed neural controller in the cart–pole loop. As in the classical benchmarks, the lines represent bounds over the entire requested disturbance interval, whereas BaB markers certify only the queried budgets. The guarantees refer to the verified learned-model dynamics. Extending either certificate to a physical plant requires a separately verified enclosure of modeling error.

Table 4: Normalized certified area A for learning-enabled systems. The guarantee applies to the modeled vector field. The last column gives GAIN-Cert runtime in seconds.

Learned setting	Staircase	Direct affine-level	Uniform neural	GAIN-Cert	GAIN-Cert time (s)
Learned FHN dynamics	0.416	0.398	0.394	**0.350**	77.0
Cart-pole + neural control	0.461	0.440	0.435	**0.391**	199.1

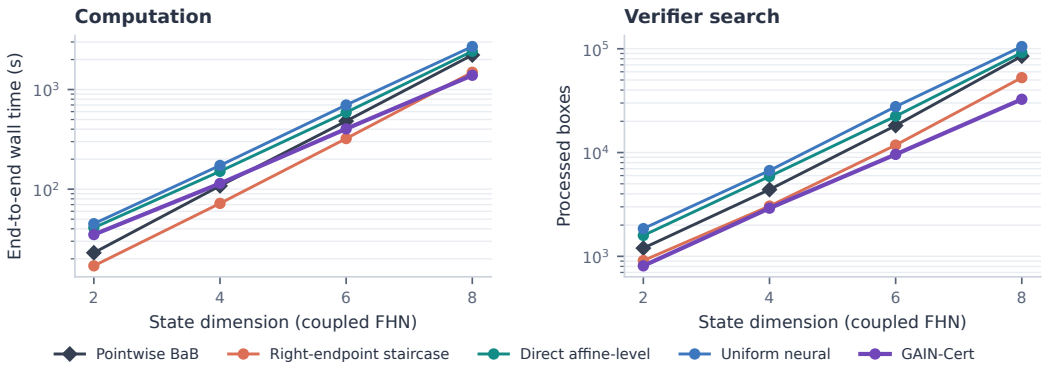


Figure 5: Coupled FitzHugh–Nagumo dimension stress test. Both time and processed boxes are recorded on logarithmic axes.

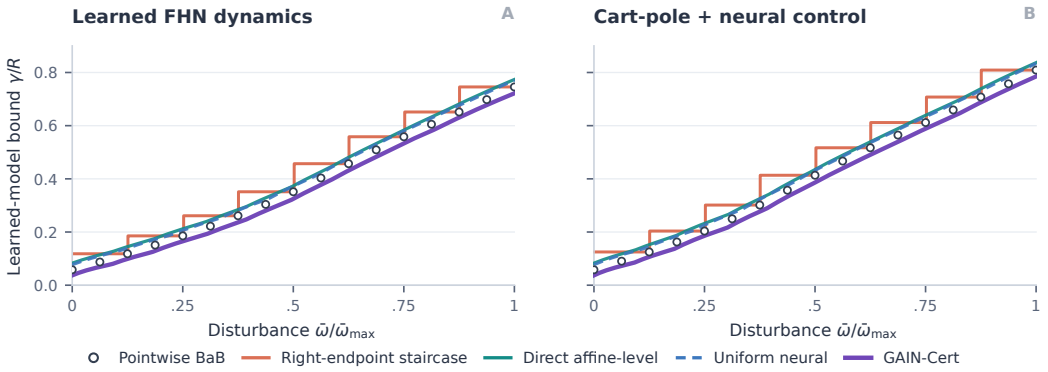


Figure 6: Learning-enabled examples: a neural ODE fitted to an excitable neuron and a cart–pole with a fixed neural feedback controller.

Table 4 quantifies the full-range curves in Fig. 6. In the reported results, GAIN-Cert has the smallest normalized area on both learned settings: 0.350 versus 0.394 for the next-smallest curve on learned FitzHugh–Nagumo, and 0.391 versus 0.435 for the cart–pole with neural feedback. Its end-to-end runtimes are 77.0 and 199.1 seconds, respectively. Areas compare tightness *within* each setting, whose normalization and disturbance range need not agree with the other setting.

Polynomial comparator. Table 5 adds a scoped SOS sweep on the four polynomial vector fields; the nonpolynomial pendulum and cart–pole models are omitted. Both methods use the same fixed V, dissipation target, region, and disturbance budgets, so the area comparison concerns how a curve is certified under those shared choices.

Table 5: Scoped polynomial SOS comparison, using the same fixed V, κ, operating region, and disturbance budgets. ΔA is the SOS sweep area minus GAIN-Cert area. Nonpolynomial dynamics are omitted.

Polynomial benchmark	SOS sweep	GAIN-Cert	ΔA
Stable linear	0.413	**0.369**	+0.044
Bistable Duffing	0.426	**0.368**	+0.058
FitzHugh–Nagumo	0.382	**0.339**	+0.043
Stable Lorenz	0.407	**0.352**	+0.055

The resulting ΔA is positive on all four listed systems (0.043–0.058), indicating smaller normalized areas for GAIN-Cert than for the displayed SOS sweep. This comparison is limited to the specified SOS formulation and settings; it does not establish a general ordering between verification frameworks.