

Enforcing Safety on Aging Systems: Sub-Manifold Neural-Operator Control Under Nonlinear Degradation

Sehaj Singh

Abstract—Classical optimal control is analyzable at a fixed design point but carries no guarantee once the plant drifts from it, and real machines drift: winding temperature rises, structural wear accumulates, and torque-transfer efficiency decays nonlinearly with both. Learned controllers adapt across the resulting high-dimensional state but are normally deployed without a stability certificate. We study an architecture that keeps a learned controller inside a per-step control-Lyapunov certificate on a 100-state, single-input plant with temperature- and wear-dependent efficiency. The construction rests on a structural observation: wear is monotone along every trajectory, so no Lyapunov function positive definite in wear can decrease, and the certificate must be restricted to the regulated sub-manifold with the monotone coordinates admitted only as parameters. A partially input-convex network supplies the certificate, convex in the 98 regulated coordinates and reshaped by the 2 monotone coordinates through a nonnegative rank-one gating path; because the regulated dynamics are exactly affine in the scalar input, the min-norm repair of the learned command onto the decay half-space is analytic and solver-free. Evaluated over a 15-seed sweep on a multi-rate plant (200 Hz mechanical, 10 Hz thermal/wear) with actuator saturation, transport delay, an ADC dead-band, and split sensor noise, the certified controller matches or beats a tuned LQR on tracking in every regime, is 2.1% ahead of it under thermal/wear stress, and cuts peak overshoot by 5.3 to 7.1 times under load, with zero divergences across 45 episodes. The measured cost is 14% more control energy in the degraded regime, which we report rather than omit.

Index Terms—Lyapunov methods, neural networks, machine learning, stability of nonlinear systems, sampled-data control, fault tolerance.

I. INTRODUCTION

CONTROL of an aging actuator forces a choice between two failure modes. A fixed classical law, a PID or an LQR gain synthesized at a nominal operating point, has behavior one can analyze and bound at that point, but it carries no guarantee once the plant leaves it. In the plant studied here torque-transfer efficiency decays as

$$\eta(x) = \exp(-k_T(T - T_{\text{amb}})_+ - k_W W), \quad (1)$$

so a gain computed at $\eta = 1$ under-actuates several-fold once the machine is hot and worn. The guarantee that made it trustworthy was a guarantee about a plant that no longer exists. A learned controller can in principle adapt across the full high-dimensional state, but it is normally deployed without

a certificate that the closed loop is stable, and with no way to observe it drift toward a boundary until it has crossed one [1].

A. Contribution and scope

We study one architecture that keeps a learned controller inside a per-step control-Lyapunov certificate. A learned policy proposes a command and a closed-form projection repairs it, when needed, onto a decreasing-energy half-space. The scope is stated up front. This is one single-input simulated plant; its wear law is a modeling choice, not measured hardware; and the certificate is a guarantee about a 98-coordinate regulated subsystem, with the two monotone aging coordinates handled separately and audited rather than certified (Section IV). We do not claim a resolution of the general adaptation-safety problem. The contributions are:

- 1) A structural argument (Section IV) that monotone degradation coordinates *must* be excluded from the convex argument of any decreasing certificate, together with a construction that admits them as parameters through a nonnegative rank-one gating path, so the certified energy bowl reshapes as the machine ages.
- 2) An analytic, solver-free safety projection exploiting exact input-affinity of the regulated dynamics, with a scale-aware regularizer that bounds the repair by the actuator range (Section V).
- 3) A practical-stability relaxation with an exact ultimate-boundedness guarantee (Proposition 1), and an honest account of the conditions under which it applies.
- 4) A measured evaluation on a multi-rate sampled-data plant with hardware artifacts, reporting the regime where the method costs more than the baseline (Section VII).

B. Relation to prior work

The verification methodology is inherited from work certifying the Lyapunov conditions of Cui and Zhang’s neural controller [2] on a lossy Kron-reduced NE39 swing model using CROWN [3] and automatic bound propagation [4] over a region rather than over samples. The governing result there, that 99.98% sampled satisfaction of a decrease condition still admits genuine counterexamples a directed search finds and 2×10^5 random samples miss, is why no claim here rests on sampled satisfaction. The certificate parameterization follows the input-convex network literature [5] and its use for stable dynamics [6], [7]; the repair step is the scalar specialization of a control-barrier quadratic program [8]; the operator is a

S. Singh is an independent researcher (e-mail: sehajrsinghs@gmail.com).

Code, data, and a full replication pipeline are available at <https://github.com/sehajr-singhs/operator>.

kernel-integral map in the neural-operator family [9], [10]; and the train-versus-deploy correction is DAgger-style aggregation [11]. What is new here is the degrading plant, and the sub-manifold restriction it forces.

II. PROBLEM FORMULATION

A. Plant

The state $x \in \mathbb{R}^{100}$ comprises angular position, angular velocity, winding temperature T , wear W , and 96 latent cross-coupling modes. It splits into a *regulated* part and a *context* part,

$$z \triangleq (\theta, \omega, \ell_1, \dots, \ell_{96}) \in \mathbb{R}^{98}, \quad c \triangleq (T, W) \in \mathbb{R}^2. \quad (2)$$

The origin $z = 0$ is an equilibrium of the regulated subsystem for every c . Temperature relaxes toward ambient and wear is monotone,

$$\dot{W} = k_w |\omega| \max(T, 0) \geq 0. \quad (3)$$

Efficiency (1) multiplies the input column, so B is state dependent. Only the velocity channel is actuated, hence the regulated dynamics are *exactly* affine in u ,

$$\dot{z} = f_z(x, 0) + g_z(x) u, \quad g_z(x) = \frac{\eta(x)}{J} e_2, \quad (4)$$

with no residual. This is what collapses the safety filter to a scalar projection in Section V.

B. Multi-rate sampled-data loop

Evaluation replaces idealized state access with a sampled-data feedback path. The mechanical subsystem is integrated by RK4 at $\Delta t = 5$ ms (200 Hz) with the context frozen at every stage; the thermal and wear channels accumulate and commit on a 10 Hz boundary, with wear ratcheted so integration error cannot walk it backward. The feedback path adds a two-tick (10 ms) transport delay, an ADC dead-band that zeros sub-LSB mechanical readings, hard torque limits $u \in [-12, 12]$ N-m, and, in the stochastic regime, high-frequency Gaussian noise on the fast channels with a slow random-walk bias on the temperature channel.

III. ARCHITECTURE

The controller composes three learned stages with one classical safety layer.

Manifold projection ($\mathbb{R}^{100} \rightarrow \mathbb{R}^{128}$): The raw state is heterogeneous, mixing radians, degrees Celsius, dimensionless wear, and latent modes. The projection whitens by physical scale, lifts the four physical channels through a random-Fourier feature map across roughly two decades of frequency, and normalizes onto a well-conditioned 128-dimensional shell.

Graph neural operator ($\mathbb{R}^{128} \rightarrow \mathbb{R}$): The embedding coordinates are treated as nodes and the command is produced by kernel-integral message passing [9],

$$v \leftarrow \sigma \left(\mathcal{W}v + \frac{1}{N} \sum_y \kappa(x, y) v(y) \right), \quad (5)$$

where κ is a multilayer perceptron of the coordinate pair, making the map discretization-aware rather than a free $N \times N$ matrix. The command is anchored at the equilibrium, $u(x) =$

$\text{net}(x) - \text{net}(x_{\text{eq}}(c))$, which forces $u = 0$ whenever $z = 0$ exactly, for every context and every parameter value. Without that anchor a residual bias of a fraction of a newton-metre parks the plant at a nonzero position where the bias balances the restoring torque.

IV. THE SUB-MANIFOLD RESTRICTION

A. The obstruction

Let V be continuously differentiable and positive definite in the full state, and suppose $\partial V / \partial W > 0$ on a set of positive measure. Along any trajectory, (3) gives $\dot{W} \geq 0$, so the contribution $(\partial V / \partial W) \dot{W}$ is nonnegative and cannot be cancelled by the input, which does not act on W . Whenever the machine is moving and warm, such a V therefore cannot satisfy $\dot{V} < 0$. No amount of training repairs this: the coordinate has to leave the convex argument. This is not a limitation of a particular parameterization but a property of monotone degradation, and it is the reason the certificate below is defined on z alone.

B. Parameterized certificate

We take $V : \mathbb{R}^{98} \times \mathbb{R}^2 \rightarrow \mathbb{R}_+$, convex in z and arbitrary in c . The context path emits the nonnegative weights themselves, as a rank-one modulation of a shared nonnegative kernel,

$$W_k^{\text{eff}}(c) = W_k^+ \odot (r_k(c) q_k(c)^\top), \quad r_k, q_k \geq 0, \quad (6)$$

so that layer k computes $h_{k+1} = \text{softplus}(W_k^{\text{eff}}(c)h_k + U_k z + b_k(c))$. The rank-one form never requires a batched matrix product, since $(W \odot (r q^\top))h = r \odot (W(q \odot h))$: gate the input, apply the shared kernel, gate the output. Convexity in z survives because $W_k^{\text{eff}} \geq 0$ elementwise, affine maps of z are convex, nonnegative combinations of convex functions are convex, and softplus is convex and nondecreasing. Positive definiteness and $V(0; c) = 0$ are then enforced structurally,

$$V(z; c) = \text{relu}(\phi(z; c) - \phi(0; c)) + \varepsilon_0 \|z\|^2, \quad (7)$$

with $\varepsilon_0 = 10^{-3}$, so $V(0; c) = 0$ exactly and $V(z; c) \geq \varepsilon_0 \|z\|^2$ for *any* parameter values. No training run can break either property.

Because c is zero-order held between slow ticks, W_k^{eff} is constant there too: the certified bowl is reshaped strictly on the 10 Hz boundary while the operator regulates the fast states at 200 Hz against a bowl that is fixed within each interval. The gating is thus the mechanism that separates high-frequency regulation from low-frequency degradation.

V. SAFETY PROJECTION AND PRACTICAL STABILITY

A. Analytic min-norm repair

Let $a \triangleq \nabla_z V^\top g_z(x)$ and $b \triangleq \nabla_z V^\top f_z(x, 0) + \alpha V$. The decay requirement is the scalar half-space $au + b \leq \epsilon$. The min-norm feasible repair of a proposed u_{raw} is

$$u^* = u_{\text{raw}} - \frac{a \text{relu}(au_{\text{raw}} + b + \delta - \epsilon)}{a^2 + \rho}, \quad (8)$$

one dot product, differentiable, with no solver in the loop. Two details are load-bearing. The strict-interior back-off $\delta =$

$10^{-6}(1 + |b|)$ prevents projection onto the exact boundary, where roughly half of all repaired steps would report as uncertified from floating-point noise alone. The regularizer is scale-aware, $\rho = \rho_0 + (|b|/2u_{\max})^2$: the repair magnitude peaks at $\text{relu}(\cdot)/(2\sqrt{\rho})$ over a , so this choice bounds it by roughly the actuator range. Without it, a state where the input is nearly orthogonal to $\nabla_z V$ produces an enormous correction that saturates the actuator, with the safety layer doing more damage than the condition it enforces.

B. Ultimate boundedness

Sensor quantization makes exact asymptotic convergence unachievable in a sampled-data loop, so we enforce and certify against $D(z) \triangleq \nabla_z V^\top f_z + \alpha V \leq \epsilon$ with $\epsilon = 10^{-4}$ rather than $D \leq 0$.

Proposition 1 (Ultimate boundedness). *Fix the context c over a slow interval. Since $\dot{V} = \nabla_z V^\top f_z$ with c held, the condition $D \leq \epsilon$ is exactly $\dot{V} \leq -\alpha V + \epsilon$. By the comparison lemma [12, Lemma 3.4],*

$$V(z(t); c) \leq e^{-\alpha t} V(z(0); c) + \frac{\epsilon}{\alpha} (1 - e^{-\alpha t}), \quad (9)$$

so $\limsup_{t \rightarrow \infty} V \leq \epsilon/\alpha$. The sublevel set $\Omega = \{z : V(z; c) \leq \epsilon/\alpha\}$ is forward invariant and attracting, and by the structural floor in (7) it lies inside the ball $\|z\| \leq \sqrt{\epsilon/(\alpha\epsilon_0)}$.

With $\epsilon = 10^{-4}$ and $\alpha = 0.5$ the guaranteed ultimate bound is $V \leq 2 \times 10^{-4}$, and the implied radius 0.45 is the conservative worst case from the quadratic floor; measured trajectories settle far inside it.

Remark 1 (Scope, stated plainly). *Proposition 1 holds along any trajectory on which the relaxed condition is actually enforced at every step. The projection enforces it wherever it has control authority ($a \neq 0$) and the actuator is unsaturated; the certificate is met on 96.9% to 100% of steps (Section VII), so ultimate boundedness is conditional on that enforcement and is not a theorem over all reachable states. The zero-divergence outcome is an empirical observation consistent with the proposition, not a direct corollary. The bound is on z ; the monotone drift is handled by a separate audit.*

Remark 2 (Intra-interval scope and the boundary jump). *Proposition 1 is stated with c fixed, and that is the only regime in which it is claimed. Let t_k denote the slow-tick instants and c_k the context held on $[t_k, t_{k+1})$, with $T_s = 0.1$ s. Within an interval the bound applies verbatim. At a boundary the gating (6) reshapes the bowl, so the certified quantity is evaluated against a different function on either side and may step discontinuously by*

$$\Delta_k \triangleq V(z(t_{k+1}); c_{k+1}) - V(z(t_{k+1}); c_k). \quad (10)$$

Chaining (9) across intervals with $\lambda = e^{-\alpha T_s}$ gives $V_{k+1} \leq \lambda V_k + (\epsilon/\alpha)(1 - \lambda) + \Delta_k$, so an episode-level bound requires a uniform jump bound $\Delta_k \leq \bar{\Delta}$, under which

$$\limsup_{k \rightarrow \infty} V_k \leq \frac{\epsilon}{\alpha} + \frac{\bar{\Delta}}{1 - \lambda}. \quad (11)$$

We do not prove such a $\bar{\Delta}$ here, and the arithmetic is unforgiving: with $\alpha = 0.5$ and $T_s = 0.1$ s, $1 - \lambda \approx 0.049$, so a per-tick jump enters (11) amplified by a factor of 20.5. We do, however, measure Δ_k directly. At each slow boundary the evaluator evaluates the certificate twice at the same z , once under c_k and once under c_{k+1} , which isolates the jump exactly rather than inferring it; Table IV reports the result over the full sweep. The measured jumps are small, and their sign structure follows the physics: the nominal mean is negative, the reshaped bowl reading slightly lower at fixed z when the machine is not degrading, and turns positive only in the two regimes where temperature and wear are actually moving. Substituting the worst observed positive jump into (11) gives an implied $\limsup V$ between 3.4×10^{-4} and 5.0×10^{-3} , that is, between 1.7 and 25 times the intra-interval bound $\epsilon/\alpha = 2 \times 10^{-4}$, and still small in absolute terms.

This is a measurement over tested trajectories, not a uniform bound over reachable states, and (11) inherits that status: it is an empirically instantiated bound, not a theorem. Establishing $\bar{\Delta}$ properly, by a Lipschitz argument on the gate in c combined with the bounded per-tick context increment, or by a dwell-time or piecewise-Lyapunov treatment of the switched system, remains the natural next theoretical step.

C. Full-state audit

The certified quantity holds c frozen. Along the real trajectory c also moves, so the honest rate is $\dot{V} = \nabla_z V^\top f_z + \nabla_c V^\top \dot{c}$. We compute $\nabla_c V$ through the same network and read $\dot{c}$ from the plant, and report the resulting full-state certification rate alongside the reduced one. Measuring this term quantifies what the sub-manifold restriction costs instead of asserting it is zero.

VI. TRAINING

The loss combines a differentiable closed-loop rollout, a decaying rehearsal against the LQR gain, a violation penalty applied to the *unprojected* command, and a physics-informed shaping term pulling V toward the plant's true mechanical-plus-latent storage function. The violation penalty is normalized by V : the unnormalized form has a degenerate minimizer, since both terms scale with V and the cheapest way to drive the penalty to zero is to shrink V onto its quadratic floor, whose gradient points along z where a single actuator has almost no authority.

The relaxation exposed a train-versus-deploy gap. With sub- ϵ violations no longer suppressed, an earlier policy over-actuated in the nominal regime, whose closed-loop trajectories visit states the offline pool under-represents. Training the operator on its own deployment states [11], every 10 iterations, closes it: nominal control effort falls from an over-actuating 182 to 22, nominal RMSE from +171% relative to LQR to +2.7%, and the nominal projection-intervention rate to zero.

VII. RESULTS

All figures are mean $\pm$ one standard deviation over 15 seeds, 10 s episodes on the multi-rate plant. Scenario A is cold and unworn; B is 95°C at wear 0.55, where efficiency falls to

TABLE I
TRACKING RMSE (RAD), 15 SEEDS

Scen.	PID	LQR	Operator
A	0.1274 ± 0.0065	0.1166 ± 0.0066	0.1198 ± 0.0067
B	0.1959 ± 0.0106	0.1527 ± 0.0085	0.1496 ± 0.0079
C	0.1554 ± 0.0084	0.1297 ± 0.0076	0.1397 ± 0.0079

TABLE II
PEAK OVERSHOOT (RAD) AND CONTROL EFFORT $\int u^2 dt$

Scen.	Overshoot		Effort	
	LQR	Operator	LQR	Operator
A	0.0007	0.0013	25.29	22.01
B	0.1842	0.0258	99.18	112.86
C	0.0524	0.0099	56.68	32.44

roughly a fifth of nominal; C is warm and worn with sensor and actuation noise.

Tracking is competitive with LQR in every regime: within 2.7% nominally, 2.1% ahead under thermal/wear stress, and 7.7% behind under stochastic disturbance, all within a few times the seed spread. Overshoot is cut by 7.1 times in the degraded regime and 5.3 times under noise, which is the axis that drives mechanical fatigue; nominally both controllers are already near zero and LQR is marginally tighter, a difference inside the operator’s own seed spread. Control energy is lower nominally and under noise but 14% higher under thermal/wear. That is the one cost we state without softening: the operator holds a hot, worn plant on a tighter, lower-overshoot trajectory and pays for it in energy. No episode diverged in any regime.

VIII. LIMITATIONS

The degraded-regime energy cost above is real and regime-dependent. Certification is high but not complete and its threshold is relaxed; the stochastic regime is lowest because the certificate is evaluated at a noise-corrupted measured state for the commanded torque, so those percentages describe the certificate arithmetic, not a physical guarantee. The guarantee is on the 98-state regulated subsystem, with the monotone coordinates covered by a measured rather than proven-zero audit that is empirical over tested trajectories. The learned terminal cost used in training is a one-step temporal-difference estimate, not a verified value function. Finally, this is a single-input simulated plant whose wear law is a modeling choice with no hardware grounding.

The sharpest theoretical limit is the one in Remark 2. Ultimate boundedness is proved *intra-interval*, with the context held, and the multi-rate loop reshapes the certificate at every 10 Hz boundary. We therefore claim a bound that holds on each slow interval, not an episode-level bound, and the missing ingredient is explicit: a uniform bound $\bar{\Delta}$ on the boundary jump (10), which (11) would then convert into an episode-level result at an amplification of roughly $1/(1-e^{-\alpha T_s}) \approx 20$ under the present gains. Two consequences are worth stating. First, the guarantee is weaker than a reader skimming Proposition 1 in isolation might assume, and the switched structure that

TABLE III
CERTIFICATE AND ROBUSTNESS (OPERATOR)

Scen.	Reduced	Full-state	Projection	Diverged
A	100.0%	100.0%	0.0%	0/15
B	99.9%	98.5%	3.4%	0/15
C	96.9%	94.7%	7.0%	0/15

TABLE IV
CERTIFICATE BOUNDARY JUMP Δ_k AT THE 10 HZ CONTEXT SWITCH, MEASURED OVER ALL SEEDS AND BOUNDARIES

Scen.	mean $ \Delta_k $	worst Δ_k^+	implied $\limsup V$
A	1.1×10^{-6}	7.0×10^{-6}	3.4×10^{-4}
B	1.3×10^{-5}	2.4×10^{-4}	5.0×10^{-3}
C	9.5×10^{-6}	1.9×10^{-4}	4.0×10^{-3}

makes the certificate adapt to aging is precisely what costs the global claim. Second, and this is the part that remains genuinely open, the jump is *measured* but not *bounded*: Table IV reports Δ_k observed directly at every context switch over all seeds, so (11) can be instantiated empirically, but a maximum over tested trajectories is not a uniform bound over reachable states. This is the same distinction that motivates regional verification over sampling in the first place, and we hold ourselves to it here: the episode-level number in Table IV is evidence, not proof. Deriving $\bar{\Delta}$ from the gate’s Lipschitz constant in c against the bounded per-tick increment of (T, W) would close it, and is the first thing we would prove. We prefer to name this boundary than to obscure it, since a switched certificate whose global behavior is asserted rather than shown is exactly the failure mode this line of work exists to avoid.

IX. CONCLUSION

Monotone degradation forbids a decreasing certificate in the full state, and the correct response is to restrict the certificate to the regulated sub-manifold and admit the aging coordinates as parameters that reshape it. Doing so yields a learned controller competitive with a tuned LQR on tracking, far gentler on the actuator under load, and certified across the degradation that invalidates the linear design, at a measured cost in degraded-regime control energy. The natural next step is to replace the per-step numerical check with a regional proof by carrying the certified-training machinery of the swing-model work onto this plant, and then to close the remaining gap on a motor test bench with real sensing and a measured degradation trajectory.

REPRODUCIBILITY

A single command, `python run_all_experiments.py`, trains from an empty checkpoint, runs the 15-seed vectorized evaluation, and regenerates every table and figure cited here. Evaluation is deterministic: repeating the scenario-A sweep reproduces the reported RMSE to full double precision.

REFERENCES

- [1] C. Dawson, S. Gao, and C. Fan, “Safe control with learned certificates: A survey of neural lyapunov, barrier, and contraction methods for robotics and control,” *IEEE Transactions on Robotics*, vol. 39, no. 3, pp. 1749–1767, 2023.
- [2] W. Cui and B. Zhang, “Lyapunov-regularized reinforcement learning for power system transient stability,” *IEEE Control Systems Letters*, vol. 6, pp. 974–979, 2022.
- [3] H. Zhang, T.-W. Weng, P.-Y. Chen, C.-J. Hsieh, and L. Daniel, “Efficient neural network robustness certification with general activation functions,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [4] K. Xu, Z. Shi, H. Zhang, Y. Wang, K.-W. Chang, M. Huang, B. Kaikhura, X. Lin, and C.-J. Hsieh, “Automatic perturbation analysis for scalable certified robustness and beyond,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [5] B. Amos, L. Xu, and J. Z. Kolter, “Input convex neural networks,” in *International Conference on Machine Learning (ICML)*, 2017.
- [6] J. Z. Kolter and G. Manek, “Learning stable deep dynamics models,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [7] Y.-C. Chang, N. Roohi, and S. Gao, “Neural lyapunov control,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [8] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, “Control barrier function based quadratic programs for safety critical systems,” *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [9] Z. Li, N. Kovachki, K. Azizzadenesheli, B. Liu, K. Bhattacharya, A. Stuart, and A. Anandkumar, “Neural operator: Graph kernel network for partial differential equations,” *arXiv preprint arXiv:2003.03485*, 2020.
- [10] —, “Fourier neural operator for parametric partial differential equations,” in *International Conference on Learning Representations (ICLR)*, 2021.
- [11] S. Ross, G. J. Gordon, and J. A. Bagnell, “A reduction of imitation learning and structured prediction to no-regret online learning,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2011.
- [12] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Prentice Hall, 2002.