Making Koopman Control Guarantees Hold on Robots

Abulikemu Abuduweili Changliu Liu
Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA
{abulikea, cliu6}@andrew.cmu.edu

Abstract

A Koopman controller can stabilize its learned latent-space linear model yet fail to stabilize the nonlinear plant. The learned model propagates predictions in latent space, while the deployed controller repeatedly re-encodes the measured physical state. We show that nominal closed-loop stability transfers to the nonlinear plant when the learned and true closed-loop dynamics agree to first order on the embedding's tangent space at the target. We derive a robustness margin for approximate matching and show that a stable open-loop latent model cannot exactly match an unstable equilibrium. We introduce Tangent-Invariant Koopman (TIK), which estimates a local plant linearization from offline transitions and enforces first-order matching to this estimate. TIK recovers the fitted LQR gain at the target, while a bilinear predictor captures state-dependent actuation during large maneuvers. Across four simulated unstable benchmarks, all 40 Standard Koopman controllers have stable latent closed loops, but their deployed nonlinear systems are locally unstable. In contrast, all 40 TIK controllers locally stabilize their respective nonlinear plants. An offline stability test using empirical curvature estimates identifies stable TIK controllers on two cart-pole systems, but is inconclusive under high noise. Under explicit model-error and curvature bounds, an invariant terminal set guarantees constraint satisfaction after entry under local feedback. TIK achieves 88% constrained cart-pole swing-up success with a 0.25 ms planning budget and controls simulated seven-joint arms within 1–5 ms per step.

Code and project page: https://github.com/Walleclipse/tik-koopman

Keywords: Koopman operator, model predictive control, stability certificates

1 Introduction

Koopman models approximate nonlinear dynamics with linear dynamics in a learned latent space, enabling efficient model-based control (Williams et al., 2015). However, a controller designed on these learned dynamics may fail to stabilize the nonlinear plant. The learned model propagates its own predictions in latent space, while the deployed controller repeatedly re-encodes the measured physical state (Figure 1a,b). Consequently, a stable latent closed loop does not guarantee stability of the deployed nonlinear system. In the double cart-pole example in Figure 1c, Standard Koopman predicts recovery from a 0.6° tilt, while the simulated plant falls.

Koopman models lift nonlinear dynamics into a higher-dimensional space, using fixed dictionaries or learned observables (Lusch et al., 2018, Shi and Meng, 2022). With linear lifted dynamics, quadratic costs, and linear constraints, model predictive control (MPC) is a convex quadratic program (QP) (Korda and Mezić, 2018). Near a target inside the constraints, a Riccati terminal cost makes its feedback law coincide with lifted linear-quadratic regulation

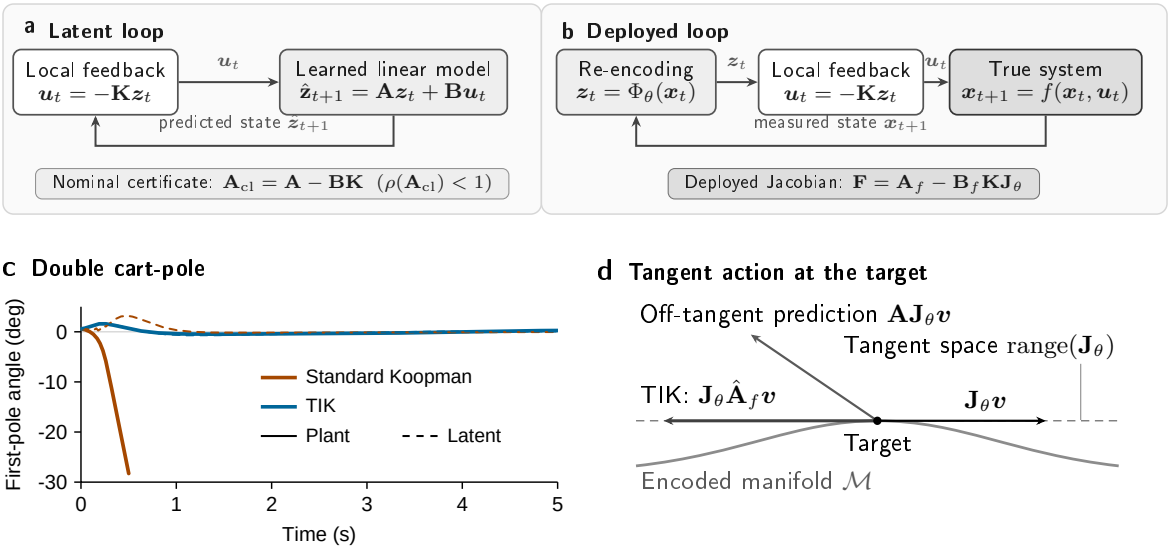


Figure 1: **The nominal and deployed loops can disagree.** (a) The learned model propagates its own predictions. (b) Deployment feeds back the re-encoded plant state. (c) First-pole angle for Standard Koopman and TIK on shared axes (seed 42, initial tilt 0.6°). Each controller has a latent prediction (dashed) and a simulated plant rollout (solid). The Standard Koopman plant trace ends when the pole falls. (d) The true linearized image of a tangent vector lies in range$(\mathbf{J}_\theta)$. A learned update can leave this space. TIK sets $\mathbf{A}\mathbf{J}_\theta\boldsymbol{v} = \mathbf{J}_\theta\hat{\mathbf{A}}_f\boldsymbol{v}$, which equals the true tangent action when the fit is exact.

(LQR). The nominal stability test checks whether the lifted closed-loop matrix $\mathbf{A}_{\mathrm{cl}}$ is Schur, with all eigenvalues inside the unit circle. Local plant stability depends on the physical Jacobian $\mathbf{F}$, which also involves the encoder's derivatives. A model with low prediction error can still give the wrong stability verdict.

What is required for a nominal stability certificate to remain valid after re-encoding? The answer lies in the tangent space of the learned embedding. If the latent closed-loop dynamics reproduce the physical closed-loop Jacobian on this space, every physical closed-loop eigenvalue also appears in the latent model (Theorem 2). We also give a margin for approximate matching, and we show that a latent model that is open-loop stable cannot match an open-loop unstable equilibrium (Corollary 3).

We then build the matching into the model. *Tangent-Invariant Koopman* (TIK) enforces the two first-order equations $\mathbf{A}\mathbf{J}_\theta = \mathbf{J}_\theta\hat{\mathbf{A}}_f$ and $\mathbf{B} = \mathbf{J}_\theta\hat{\mathbf{B}}_f$, where $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ is a local plant linearization fitted from offline transitions (Proposition 4). For fixed $\mathbf{J}_\theta$, these equations are linear in $\mathbf{A}, \mathbf{B}$. The fit determines the local feedback: TIK's gain equals the LQR gain on the fit. The learned predictor models dynamics away from the target, including state-dependent actuation during large maneuvers. We also give a stability test that accounts for estimation error (Theorem 6) and an invariant terminal set under explicit curvature and noise assumptions (Theorem 8).

In four simulated open-loop unstable benchmarks, all 40 Standard Koopman controllers pass the nominal test but have unstable physical closed-loop Jacobians; all 40 TIK controllers have Schur physical Jacobians. Further experiments evaluate offline stability verification, constrained maneuvers, and control under short planning budgets.

2 Related Work

Learning Koopman models for control. Williams et al. (2015) fit a linear operator

on a fixed dictionary. Neural methods learn the observables, often with a latent-linearity loss that penalizes errors in the lifted coordinates (Lusch et al., 2018, Shi and Meng, 2022). Linear predictors with control inputs lead to QP-based MPC (Korda and Mezić, 2018, Proctor et al., 2018), and bilinear predictors allow the input effect to depend on the state (Bruder et al., 2021). Applications include manipulation and legged locomotion (Chen et al., 2024, Jung et al., 2026, Li et al., 2025). Exact finite-dimensional embeddings have structural limits (Brunton et al., 2016, Liu et al., 2025). Van Goor et al. study reprojection onto the encoded manifold (van Goor et al., 2023), and Ge et al. (2026) study how the action operating point affects learned world-model accuracy for humanoid tracking. We study how re-encoding the measurement affects closed-loop stability at a target.

Local nonlinear stability and Koopman guarantees. Lyapunov's indirect method linearizes the closed loop in the physical state space (Elaydi, 2005, Khalil, 2002). Koopman eigenfunction theory relates physical Jacobians to stability under exact embeddings (Mauroy and Mezić, 2016). For learned predictors, Shang et al. (2026a) establish local exponential stability under a sufficiently small proportional one-step prediction error and identify vanishing first-order prediction error as a sufficient condition. SafEDMD uses proportional surrogate-error bounds for robust feedback synthesis (Nüske et al., 2023, Strässer et al., 2025, 2026a); Strässer et al. (2026b) survey how model-error bounds support closed-loop guarantees. Jain (2026) estimates the Jacobian of a candidate closed loop from near-equilibrium perturbations and uses it to construct a region of attraction. We derive an explicit singular-value margin for the closed-loop error Jacobian. TIK matches derivatives to a fitted state-input model and tests the remaining fit error in the direction induced by the deployed feedback. The resulting derivative constraints allow learned prediction and classical local stability analysis to be used together. Open-loop stable parameterizations (Fan et al., 2026, Mamakoukas et al., 2023) encode useful prior knowledge for stable plants. Our obstruction concerns using them at unstable equilibria.

3 Background and Problem Setup

We consider a discrete-time nonlinear system $\boldsymbol{x}_{t+1} = f(\boldsymbol{x}_t, \boldsymbol{u}_t)$ with physical state $\boldsymbol{x}_t \in \mathbb{R}^{n_x}$, input $\boldsymbol{u}_t \in \mathbb{R}^{n_u}$, and unknown continuously differentiable dynamics f. The learner has logged transitions $(\boldsymbol{x}_t, \boldsymbol{u}_t, \boldsymbol{x}_{t+1})$ and a specified equilibrium $(\boldsymbol{x}^*, \boldsymbol{u}^*)$ with $f(\boldsymbol{x}^*, \boldsymbol{u}^*) = \boldsymbol{x}^*$. The goal is to stabilize this equilibrium under state and input limits. We shift coordinates so that $(\boldsymbol{x}^*, \boldsymbol{u}^*) = (\mathbf{0}, \mathbf{0})$ (Appendix C keeps $(\boldsymbol{x}^*, \boldsymbol{u}^*)$ explicit) and write $\mathbf{A}_f = \partial f/\partial \boldsymbol{x}|_{(\mathbf{0},\mathbf{0})}$ and $\mathbf{B}_f = \partial f/\partial \boldsymbol{u}|_{(\mathbf{0},\mathbf{0})}$ for the true Jacobians at the target. Let $\rho(\cdot)$ denote the spectral radius (largest eigenvalue magnitude). The equilibrium is *open-loop unstable* if $\rho(\mathbf{A}_f) > 1$, as in an upright pendulum or cart-pole; stabilizing such systems requires active feedback and is our focus.

Neural Koopman model. The Koopman operator advances observables by composition with the dynamics. For a fixed input $\boldsymbol{u}$, this action is $\psi \mapsto \psi(f(\cdot, \boldsymbol{u}))$ (Koopman, 1931, Korda and Mezić, 2018, Proctor et al., 2018). It is linear even for nonlinear f, but acts on an infinite-dimensional function space. In practice, it is approximated on N observables stacked into an embedding $\Phi_\theta : \mathbb{R}^{n_x} \to \mathbb{R}^N$ that concatenates $\boldsymbol{x}$ with a multilayer perceptron (MLP) encoder $\Psi_\theta : \mathbb{R}^{n_x} \to \mathbb{R}^{N-n_x}$ (writing $[\boldsymbol{a}; \boldsymbol{b}]$ for column stacking) (Chen et al., 2024, Jung et al., 2026, Li et al., 2025, Shi and Meng, 2022):

$$\boldsymbol{z}_t = \Phi_\theta(\boldsymbol{x}_t) = \big[\boldsymbol{x}_t;\, \Psi_\theta(\boldsymbol{x}_t) - \Psi_\theta(\mathbf{0})\big], \qquad \hat{\boldsymbol{z}}_{t+1} = \mathbf{A}\boldsymbol{z}_t + \mathbf{B}\boldsymbol{u}_t, \qquad \boldsymbol{x}_t = \mathbf{C}\boldsymbol{z}_t, \quad \mathbf{C} = [\mathbf{I}_{n_x} \;\; \mathbf{0}]. \tag{1}$$

Thus $\Phi_\theta(\mathbf{0}) = \mathbf{0}$ and $\mathbf{C}$ reads out the state linearly, keeping state costs and bounds convex in $\boldsymbol{z}$. Because finite observable sets rarely span an exact invariant subspace (Brunton et al., 2016), $\hat{\boldsymbol{z}}_{t+1}$ generally differs from $\Phi_\theta(\boldsymbol{x}_{t+1})$. *Standard Koopman* trains $(\Phi_\theta, \mathbf{A}, \mathbf{B})$ on offline trajectories with a multi-step state-prediction loss and an inverse control loss (Abuduweili et al., 2026); *Koopman + invariance penalty* adds a latent-linearity loss following Lusch et al. (2018) (Appendix E.1). This loss penalizes feature prediction errors on training trajectories. Tangent invariance concerns derivatives at the target.

Koopman MPC and LQR. Given $\boldsymbol{z}_0 = \Phi_\theta(\boldsymbol{x}_t)$, Koopman MPC solves the convex QP

$$\begin{aligned} \min_{\boldsymbol{u}_{0:N_h-1}} \ & \textstyle\sum_{k=0}^{N_h-1} \big(\boldsymbol{z}_k^\mathsf{T}\mathbf{C}^\mathsf{T}\mathbf{Q}_x\mathbf{C}\boldsymbol{z}_k + \boldsymbol{u}_k^\mathsf{T}\mathbf{R}\boldsymbol{u}_k\big) + \boldsymbol{z}_{N_h}^\mathsf{T}\mathbf{P}\boldsymbol{z}_{N_h} \\ \text{s.t.} \ & \boldsymbol{z}_{k+1} = \mathbf{A}\boldsymbol{z}_k + \mathbf{B}\boldsymbol{u}_k, \quad \boldsymbol{x}_{\min} \le \mathbf{C}\boldsymbol{z}_k \le \boldsymbol{x}_{\max}, \quad |\boldsymbol{u}_k| \le \boldsymbol{u}_{\max}, \end{aligned} \tag{2}$$

and applies $\boldsymbol{u}_t = \boldsymbol{u}_0$ (Korda and Mezić, 2018), where $\mathbf{Q}_x \succeq 0$, $\mathbf{R} \succ 0$, and $\mathbf{P} \succeq 0$ solves the Riccati equation of $(\mathbf{A}, \mathbf{B}, \mathbf{C}^\mathsf{T}\mathbf{Q}_x\mathbf{C}, \mathbf{R})$. When constraints are inactive, (2) gives the Koopman LQR policy, $\boldsymbol{u}_t = -\mathbf{K}\boldsymbol{z}_t$ with $\mathbf{K} = (\mathbf{R} + \mathbf{B}^\mathsf{T}\mathbf{P}\mathbf{B})^{-1}\mathbf{B}^\mathsf{T}\mathbf{P}\mathbf{A}$ (Proposition 7). Both policies apply the same input in a neighborhood of the target and thus induce the same local closed-loop Jacobian.

3.1 The Latent Loop versus the Deployed Loop

Latent loop. The model feeds back its own predictions, giving $\boldsymbol{z}_{t+1} = \mathbf{A}_{\mathrm{cl}}\boldsymbol{z}_t$ with $\mathbf{A}_{\mathrm{cl}} = \mathbf{A} - \mathbf{B}\mathbf{K}$. The nominal stability certificate is $\rho(\mathbf{A}_{\mathrm{cl}}) < 1$ ($\mathbf{A}_{\mathrm{cl}}$ is Schur). With $\mathbf{R} \succ 0$, stabilizability of $(\mathbf{A}, \mathbf{B})$ and detectability of $(\mathbf{A}, \mathbf{Q}_x^{1/2}\mathbf{C})$ ensure a stabilizing LQR solution.

Deployed loop. The robot measures $\boldsymbol{x}_t$, encodes $\boldsymbol{z}_t = \Phi_\theta(\boldsymbol{x}_t)$, and applies $\boldsymbol{u}_t = -\mathbf{K}\Phi_\theta(\boldsymbol{x}_t)$ to the true plant (Figure 1b). With target embedding Jacobian $\mathbf{J}_\theta := \partial\Phi_\theta/\partial\boldsymbol{x}|_{\mathbf{0}} = [\mathbf{I}_{n_x}; \partial\Psi_\theta/\partial\boldsymbol{x}|_{\mathbf{0}}] \in \mathbb{R}^{N\times n_x}$ ($\mathbf{C}\mathbf{J}_\theta = \mathbf{I}_{n_x}$, full column rank), the deployed loop and its target Jacobian are

$$\boldsymbol{x}_{t+1} = g(\boldsymbol{x}_t) := f\big(\boldsymbol{x}_t, -\mathbf{K}\Phi_\theta(\boldsymbol{x}_t)\big), \qquad \mathbf{F} := \frac{\partial g}{\partial \boldsymbol{x}}\Big|_{\mathbf{0}} = \mathbf{A}_f - \mathbf{B}_f\mathbf{K}\mathbf{J}_\theta. \tag{3}$$

By Lyapunov's indirect method (Khalil, 2002), the target is locally exponentially stable on the true system if $\rho(\mathbf{F}) < 1$ and unstable if $\rho(\mathbf{F}) > 1$. The nominal certificate checks the $N \times N$ matrix $\mathbf{A}_{\mathrm{cl}}$, whereas the robot is governed by the $n_x \times n_x$ matrix $\mathbf{F}$. For a ReLU encoder, the Jacobian is constant in a neighborhood of a target where the activation pattern is fixed (Appendix E.3).

4 When Does the Nominal Certificate Hold on the Robot?

We now compare $\mathbf{A}_{\mathrm{cl}}$ with $\mathbf{F}$ and identify when the nominal certificate transfers to the robot (Proposition 7 extends linear-gain results to Koopman MPC near the target).

One-step latent error. Given $(\boldsymbol{x}_t, \boldsymbol{u}_t)$, the true next encoded state is $\Phi_\theta(f(\boldsymbol{x}_t, \boldsymbol{u}_t))$ and the linear model predicts $\mathbf{A}\Phi_\theta(\boldsymbol{x}_t) + \mathbf{B}\boldsymbol{u}_t$, giving the one-step latent prediction error $E(\boldsymbol{x}, \boldsymbol{u}) := \Phi_\theta(f(\boldsymbol{x}, \boldsymbol{u})) - \mathbf{A}\Phi_\theta(\boldsymbol{x}) - \mathbf{B}\boldsymbol{u}$. The physical prediction loss directly supervises the first n_x coordinates $\mathbf{C}E$. Multi-step training can influence the other coordinates through later state predictions, but it imposes no exact derivative constraints at the target. LQR uses the full latent model, including the learned feature coordinates.

Definition 1 (Tangent invariance)**.** The model is *tangent-invariant* at the target $(\mathbf{0}, \mathbf{0})$ if the first derivatives of $E(\boldsymbol{x}, \boldsymbol{u})$ vanish there:

$$\frac{\partial E}{\partial \boldsymbol{x}}\Big|_{(\mathbf{0},\mathbf{0})} = \mathbf{J}_\theta\mathbf{A}_f - \mathbf{A}\mathbf{J}_\theta = \mathbf{0}, \qquad \frac{\partial E}{\partial \boldsymbol{u}}\Big|_{(\mathbf{0},\mathbf{0})} = \mathbf{J}_\theta\mathbf{B}_f - \mathbf{B} = \mathbf{0}. \tag{4}$$

Near a smooth target, the encoded states form a manifold $\mathcal{M}$ with tangent space range($\mathbf{J}_\theta$) (Figure 1d). Tangent invariance here means exact first-order matching, not just preserving this subspace: $\mathbf{A}\mathbf{J}_\theta \boldsymbol{v} = \mathbf{J}_\theta \mathbf{A}_f \boldsymbol{v}$ and $\mathbf{B}\boldsymbol{w} = \mathbf{J}_\theta \mathbf{B}_f \boldsymbol{w}$ for state and input perturbations $\boldsymbol{v}, \boldsymbol{w}$. For fixed $\mathbf{J}_\theta$, these are linear constraints on $(\mathbf{A}, \mathbf{B})$. Along the closed loop $\boldsymbol{u} = -\mathbf{K}\Phi_\theta(\boldsymbol{x})$, the one-step latent error is $\Phi_\theta(g(\boldsymbol{x})) - \mathbf{A}_{\mathrm{cl}}\Phi_\theta(\boldsymbol{x})$, with target Jacobian

$$\mathcal{E} \;:=\; \mathbf{J}_\theta \mathbf{F} - \mathbf{A}_{\mathrm{cl}} \mathbf{J}_\theta \;=\; (\mathbf{J}_\theta \mathbf{A}_f - \mathbf{A}\mathbf{J}_\theta) - (\mathbf{J}_\theta \mathbf{B}_f - \mathbf{B})\mathbf{K}\mathbf{J}_\theta, \tag{5}$$

Exact open-loop matching makes $\mathcal{E} = \mathbf{0}$ for every gain. For a fixed gain, the two errors in (5) can cancel even when each is nonzero. For a Schur matrix $\mathbf{M}$, let $d_1(\mathbf{M}) = \min_{|s|=1} \sigma_{\min}(s\mathbf{I} - \mathbf{M})$ be its distance to the nearest complex matrix with a unit-circle eigenvalue.

Theorem 2 (Transfer of the nominal stability certificate)**.** *For the state-retaining lifting in (1), assume differentiability at the target and inactive input limits nearby. Since $\mathbf{C}\mathbf{J}_\theta = \mathbf{I}$ and $\|\mathbf{C}\|_2 = 1$, each eigenvalue λ of $\mathbf{F}$ satisfies $\sigma_{\min}(\lambda\mathbf{I} - \mathbf{A}_{\mathrm{cl}}) \le \|\mathcal{E}\|_2$. Consequently: (a) if $\mathcal{E} = \mathbf{0}$ for the deployed gain, every eigenvalue of $\mathbf{F}$ is an eigenvalue of $\mathbf{A}_{\mathrm{cl}}$, so $\rho(\mathbf{F}) \le \rho(\mathbf{A}_{\mathrm{cl}})$ and $\rho(\mathbf{A}_{\mathrm{cl}}) < 1$ guarantees local exponential stability on the true system; (b) if $\mathbf{A}_{\mathrm{cl}}$ is Schur and $\|\mathcal{E}\|_2 < d_1(\mathbf{A}_{\mathrm{cl}})$, the target equilibrium is locally exponentially stable.*

Exact matching embeds every physical closed-loop eigenvalue in the latent model. Approximate matching preserves stability when the derivative error is smaller than the complex stability radius $d_1(\mathbf{A}_{\mathrm{cl}})$. This margin accounts for sensitivity to perturbations, including non-normal dynamics. It is sufficient and may be conservative; failing the test does not imply instability. Proofs are in Appendix C.

Corollary 3 (Obstruction at unstable equilibria)**.** *At an equilibrium with $\rho(\mathbf{A}_f) > 1$, tangent invariance implies $\rho(\mathbf{A}) > 1$.*

Exact tangent matching embeds every eigenvalue of $\mathbf{A}_f$ in $\mathbf{A}$. A Schur $\mathbf{A}$ therefore cannot satisfy these conditions at an unstable target. This does not rule out its use as an approximate predictor. Appendix B illustrates how exact one-step physical prediction can coexist with a misleading latent certificate.

5 Tangent-Invariant Koopman (TIK)

TIK uses near-target transitions to estimate the local plant dynamics and trains a lifted model for prediction beyond that neighborhood. Algorithm 1 summarizes the method.

Algorithm 1: **Tangent-Invariant Koopman (TIK)**

Inputs. Offline trajectories, target equilibrium, costs $\mathbf{Q}_x, \mathbf{R}$, and a Standard Koopman checkpoint.

1. Fit $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ using (6); require a full-rank weighted design.
2. Initialize the encoder and $\mathbf{W}$ from the checkpoint and any bilinear terms at zero. Fine-tune with (7) rebuilt at every update. Require the LQR conditions in Proposition 4(b); stop if they fail.
3. Solve the latent Riccati equation, set $\hat{\mathbf{K}} = \mathbf{K}\mathbf{J}_\theta$, and form $\hat{\mathbf{F}} = \hat{\mathbf{A}}_f - \hat{\mathbf{B}}_f \hat{\mathbf{K}}$.
4. With curvature and noise bounds, test the stability margin (Theorem 6). For dual-mode control, compute an invariant Ω_c if Theorem 8 holds. Otherwise report the corresponding guarantee as inconclusive.
5. For local regulation, solve the linear-model QP. For large maneuvers, plan with the bilinear predictor. If a terminal set was obtained, switch to $\boldsymbol{u} = -\hat{\mathbf{K}}\boldsymbol{x}$ on entry to Ω_c.

5.1 Fitting and Parameterizing the Lifted Dynamics

Local plant fit. The first-order Taylor model of f describes the dynamics near the target. The plant derivatives are unknown, so we estimate them from M offline transitions near the target. For $\boldsymbol{p}_i = [\boldsymbol{x}_i; \boldsymbol{u}_i]$, solve

$$(\hat{\mathbf{D}}, \hat{\boldsymbol{c}}) = \arg\min_{\mathbf{D},\boldsymbol{c}} \textstyle\sum_{i=1}^{M} w_i \|\boldsymbol{x}_{i+1} - \mathbf{D}\boldsymbol{p}_i - \boldsymbol{c}\|_2^2, \qquad \hat{\mathbf{D}} = [\hat{\mathbf{A}}_f \;\; \hat{\mathbf{B}}_f]. \tag{6}$$

Positive weights decrease with distance to the target. The intercept accounts for an offset in the fit. The centered lifted model still fixes the target at zero.

Tangent-invariant parameterization. Let $\mathbf{P}_\perp = \mathbf{I} - \mathbf{J}_\theta \mathbf{J}_\theta^+$ be the orthogonal projector onto the complement of the tangent space, where $\mathbf{J}_\theta^+$ is the Moore–Penrose pseudoinverse. TIK sets

$$\mathbf{A}(\theta, \mathbf{W}) = \mathbf{J}_\theta \hat{\mathbf{A}}_f \mathbf{J}_\theta^+ + \mathbf{W}\mathbf{P}_\perp, \qquad \mathbf{B}(\theta) = \mathbf{J}_\theta \hat{\mathbf{B}}_f. \tag{7}$$

The fit determines $\mathbf{A}$ on the tangent space and $\mathbf{B}$ at the target. The trainable $\mathbf{W}$ determines the action of $\mathbf{A}$ on the complement. In our experiments, we initialize θ and $\mathbf{W}$ from the encoder and $\mathbf{A}$ of the same-seed Standard Koopman checkpoint. At each update, we recompute $\mathbf{J}_\theta$ and rebuild $\mathbf{A}, \mathbf{B}$. Automatic differentiation propagates prediction-loss gradients through $\mathbf{J}_\theta$ and its pseudoinverse to θ. The parameterization enforces matching to the fit at every update. A hinge penalty encourages the dynamics on the orthogonal complement to be Schur, a condition we check for the LQR guarantee (Appendix F.2).

State-dependent actuation. The linear predictor has the same input matrix at every state. In swing-up, the effect of a cart force on pole acceleration can change sign between the hanging and upright configurations. An optional bilinear term (Bruder et al., 2021) gives the predictor the input matrix $\mathbf{B} + [\mathbf{N}_1\boldsymbol{z} \;\cdots\; \mathbf{N}_{n_u}\boldsymbol{z}]$ at state $\boldsymbol{z}$. This adds $\sum_{j=1}^{n_u} u_j \mathbf{N}_j \boldsymbol{z}$ to the predicted update. The added term and its first derivatives vanish at the centered target. Section 6.3 evaluates its effect on swing-up performance.

Proposition 4 (Structural properties of TIK)**.** *For any differentiable lifting at the target, (7) satisfies: (a)* $\mathbf{A}\mathbf{J}_\theta = \mathbf{J}_\theta \hat{\mathbf{A}}_f$ *and* $\mathbf{B} = \mathbf{J}_\theta \hat{\mathbf{B}}_f$*, with* $\mathcal{E} = \mathbf{J}_\theta(\mathbf{A}_f - \hat{\mathbf{A}}_f) - \mathbf{J}_\theta(\mathbf{B}_f - \hat{\mathbf{B}}_f)\mathbf{K}\mathbf{J}_\theta$*. (b) If* $\mathbf{P}_\perp \mathbf{W} \mathbf{P}_\perp$ *is Schur,* $\mathbf{R} \succ 0$*,* $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ *is stabilizable, and* $(\hat{\mathbf{A}}_f, \mathbf{Q}_x^{1/2})$ *is detectable, then* $\mathbf{K}\mathbf{J}_\theta = \hat{\mathbf{K}}$*, the physical LQR gain for* $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f, \mathbf{Q}_x, \mathbf{R})$*. The deployed loop has Jacobian* $\mathbf{F} = \mathbf{A}_f - \mathbf{B}_f \hat{\mathbf{K}}$*. (c) The bilinear term preserves (a) and (b) for the local linearization and its LQR.*

Proposition 4 separates local stabilization from nonlinear prediction. TIK recovers the fitted physical LQR gain at the target. The learned observables and bilinear terms model dynamics away from it. Any improvement over fitted LQR must therefore occur beyond the local linear regime. The proof is in Appendix C.

5.2 A Stability Certificate from Offline Data

Stability of the fitted closed loop does not establish plant stability. We bound its derivative error under the deployed feedback and compare it with the *real stability radius* $r_{\mathbb{R}}(\hat{\mathbf{F}})$, the smallest spectral norm of a real perturbation that places an eigenvalue on the unit circle (Qiu et al., 1995). Below this radius, the fitted and true Jacobians have the same stability classification.

Assumption 5 (Local curvature and noise bounds)**.** The target and regressors $\boldsymbol{p}_i$ lie in a convex neighborhood $\mathcal{N}$ on which f is twice continuously differentiable. For $\boldsymbol{p} \in \mathcal{N}$, $\|D^2 f(\boldsymbol{p})[\boldsymbol{a}, \boldsymbol{b}]\|_2 \le L\|\boldsymbol{a}\|_2\|\boldsymbol{b}\|_2$. The regressors are exact and the measured successor obeys $\|\boldsymbol{x}_{i+1} - f(\boldsymbol{p}_i)\|_2 \le \epsilon$.

Write $\mathbf{D}_f = [\mathbf{A}_f \quad \mathbf{B}_f]$. Taylor's theorem bounds the residual $\boldsymbol{\ell}_i = \boldsymbol{x}_{i+1} - \mathbf{D}_f \boldsymbol{p}_i$ by $\varrho_i = \frac{L}{2}\|\boldsymbol{p}_i\|_2^2 + \epsilon$. Let $\boldsymbol{q}_i = [\boldsymbol{p}_i; 1]$. The weighted fit requires $\boldsymbol{\Sigma} = \sum_i w_i \boldsymbol{q}_i \boldsymbol{q}_i^\mathsf{T} \succ 0$. Let $\boldsymbol{g}_{i,D}$ be the first $n_x + n_u$ entries of $\boldsymbol{\Sigma}^{-1}\boldsymbol{q}_i$. The normal equations and the triangle inequality give

$$\|(\hat{\mathbf{D}} - \mathbf{D}_f)\boldsymbol{\Gamma}\|_2 \le \eta(\boldsymbol{\Gamma}) := \textstyle\sum_{i=1}^{M} w_i \varrho_i \|\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D}\|_2. \tag{8}$$

The bound holds for each dimensionally compatible $\boldsymbol{\Gamma}$. We use $\boldsymbol{\Gamma} = [\mathbf{I}_{n_x}; -\mathbf{K}_o]$ to bound error along the state-input perturbations induced by the deployed policy, rather than over all input directions.

Theorem 6 (Stability certificate from offline data)**.** *Suppose the deployed policy preserves the target equilibrium, is unsaturated near it, and has Jacobian $-\mathbf{K}_o$ there. For lifted LQR, $\mathbf{K}_o = \mathbf{K}\mathbf{J}_\theta$, which equals $\hat{\mathbf{K}}$ for TIK under Proposition 4(b). Let $\hat{\mathbf{F}} = \hat{\mathbf{D}}\boldsymbol{\Gamma}$ have no eigenvalue on the unit circle. Under Assumption 5 and the full-rank fit condition, if $\eta(\boldsymbol{\Gamma}) < r_{\mathbb{R}}(\hat{\mathbf{F}})$, the target is locally exponentially stable on the plant when $\rho(\hat{\mathbf{F}}) < 1$ and unstable when $\rho(\hat{\mathbf{F}}) > 1$.*

The test is inconclusive if the strict inequality cannot be established. A verified lower bound on the radius is sufficient. Formal certification requires valid curvature and noise bounds; our experimental curvature estimates are empirical. Appendix C gives the proof, a conservative numerical test, and extensions to noisy regressors and cubic remainders.

5.3 Koopman MPC and a Certified Terminal Set

Proposition 7 (Koopman MPC near the target)**.** *Consider (2) with $\mathbf{R} \succ 0$, stabilizing Riccati terminal weight $\mathbf{P}$, and the target in the strict interior of the constraints. Near the target, where no constraint is active, the MPC law equals the Koopman LQR law $\boldsymbol{u} = -\mathbf{K}\Phi_\theta(\boldsymbol{x})$ for all $N_h \ge 1$. Under Proposition 4(b), TIK's deployed MPC loop has Jacobian $\mathbf{F} = \mathbf{A}_f - \mathbf{B}_f\hat{\mathbf{K}}$, and Theorem 6 applies.*

Local regulation uses the convex linear-model QP in (2) (Appendix C). For large maneuvers, we use the bilinear predictor with model predictive path integral control (MPPI) (Williams et al., 2017) or the QP-MPC variant in Appendix F.5. On entry to Ω_c, the controller switches to $\boldsymbol{u} = -\hat{\mathbf{K}}\boldsymbol{x}$.

Theorem 8 (Certified invariant terminal set)**.** *Let $\hat{\mathbf{F}}$ be Schur, $\mathbf{S} \succ 0$ solve $\hat{\mathbf{F}}^\mathsf{T}\mathbf{S}\hat{\mathbf{F}} - \mathbf{S} = -\mathbf{Q}$ for $\mathbf{Q} \succ 0$, and set $\hat{\vartheta} = \|\mathbf{S}^{1/2}\hat{\mathbf{F}}\mathbf{S}^{-1/2}\|_2 < 1$ and $\kappa = \lambda_{\max}(\mathbf{S})/\lambda_{\min}(\mathbf{S})$. Suppose $g_{\mathrm{loc}}(\boldsymbol{x}) = f(\boldsymbol{x}, -\hat{\mathbf{K}}\boldsymbol{x})$ satisfies $\|\mathbf{F} - \hat{\mathbf{F}}\|_2 \le \eta$ and $\|g_{\mathrm{loc}}(\boldsymbol{x}) - \mathbf{F}\boldsymbol{x}\|_2 \le \frac{1}{2}L_H\|\boldsymbol{x}\|_2^2 + \frac{1}{6}L_T\|\boldsymbol{x}\|_2^3$ on a convex set $\mathcal{N}$. Choose $r > 0$ such that $\hat{\vartheta} + \sqrt{\kappa}(\eta + \frac{1}{2}L_H r + \frac{1}{6}L_T r^2) \le 1$. Every ellipsoid $\Omega_c = \{\boldsymbol{x} : \boldsymbol{x}^\mathsf{T}\mathbf{S}\boldsymbol{x} \le c\}$ with $c \le \lambda_{\min}(\mathbf{S})r^2$ contained in $\mathcal{N}$ and satisfying the state and local feedback input limits is forward invariant under g_{loc}. If the bound is strict, trajectories in Ω_c converge exponentially to $\mathbf{0}$.*

After entry into Ω_c, the local controller keeps subsequent sampled states and inputs within their limits under the stated bounds. The theorem guarantees neither safe entry nor safety between sampling instants. Appendix D gives the proof and the Lyapunov-metric formulation.

6 Experiments

All experiments use simulated plants (Appendix E), and all feedback is tested on the nonlinear simulator (Appendix E.3). *Oracle* LQR and linear MPC use the simulator's target Jacobians, obtained by finite differences of its one-step map, as known-dynamics references. *Fitted* linear

Table 1: **Local plant stability at the upright target.** Plant-stable models among those with a Schur latent closed loop (Koopman rows) or among evaluated models (oracle MPC, neural NMPC); 10 seeds per learned method, one oracle per system. A dash indicates that no run yielded a stabilizing LQR gain.

	Plant-stable models (counts)			
Controller	Cart-pole (ODE)	Cart-pole (MuJoCo)	Double cart-pole	Double pendulum
Oracle linear MPC	1/1	1/1	1/1	1/1
Neural model + NMPC	10/10	10/10	1/10	0/10
Standard Koopman	0/10	0/10	0/10	0/10
Koopman + invariance penalty	10/10	10/10	0/10	7/10
EDMD (Korda and Mezić, 2018)	0/7	0/6	–	9/10
Deep Koopman with control (Shi and Meng, 2022)	10/10	10/10	0/10	9/10
Stable operator (Mamakoukas et al., 2023)	0/10	0/10	0/10	0/10
Stable embedding (Fan et al., 2026)	0/10	0/10	0/10	0/10
TIK (ours)	10/10	10/10	10/10	10/10

LQR/MPC and TIK instead estimate $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ from logged transitions (6); neural NMPC linearizes its learned predictor at each step (Appendix E.1). Simulator derivatives provide the stability reference and oracle models; controller tuning is described in the appendix. The upright Koopman MPC comparison uses $N_h = 20$ and identity LQR weights.

6.1 Does the Nominal Certificate Hold on the True System? (Q1)

All 40 Standard Koopman models have Schur latent closed loops but unstable physical Jacobians at the target. All 40 TIK models have Schur physical Jacobians (Table 1). The soft penalty stabilizes all single cart-pole models and seven double-pendulum models, but no double cart-pole model. Neither stable-by-construction variant stabilizes any of its 40 models. EDMD stabilizes nine double-pendulum models; its double cart-pole Riccati solve returned no stabilizing gain, which alone leaves controllability undetermined.

Prediction and sensitivity to feedback. Standard Koopman and the soft-penalty model have nearly identical 14-step state errors on the ODE cart-pole (0.013 and 0.012), but differ in plant stability. The neural model yields lower prediction error on the double cart-pole and double pendulum, yet only one of its 20 NMPC controllers is locally stable. To test whether nominal stability remains informative as feedback strength changes, we scale the nominal gain by 80 values $k \in [0.05, 4]$. Nominal and deployed verdicts agree for 14% of Standard Koopman model–gain pairs and 100% for TIK (Figure 2 in Appendix F.1).

Recovery from initial tilts. TIK uses the fitted physical LQR gain (Proposition 4) and matches oracle LQR's recovery outcome in 420 of 420 model–tilt tests across 40 models, versus 0 for linearized Standard Koopman (Figure 3; protocol in Appendix E.3). This comparison uses local linear feedback. For larger tilts, we start from the offline TIK checkpoint and add rollouts of its lifted policy from tilted starts to the training data. The median recovery range increased from 0.06 to 0.15 rad on double cart-pole and from 0.40 to 1.50 rad on double pendulum. All 39 completed TIK runs yield stable controllers. One of the 40 seeds diverged during training. The same training without tangent matching yields 2 stable controllers out of 30 (Appendix F.2).

Table 2: **Empirical offline stability-margin tests.** Controllers passing the estimated margin test in Theorem 6 (20 per method per row). Simulator Jacobians classify zero Standard Koopman and all 20 TIK controllers as locally stable.

			Passed estimated margin test	
System	Offline data regime	Noise ϵ	Standard Koopman	**TIK**
Cart-pole (upright)	Near target (400 pts)	0	0/20	20/20
	Two scales	0	0/20	20/20
	Two scales, noisy states	10^{-7}	0/20	20/20
	Closed-loop trajectories	3×10^{-8}	0/20	20/20
Double cart-pole (upright)	Two scales	0	0/20	20/20
	Closed-loop trajectories	10^{-10}	0/20	20/20
	Two scales (high noise)	10^{-6}	0/20	0/20

Table 3: **Dual-mode MPC for constrained cart-pole swing-up and recovery across compute budgets.** Share of rollouts (12 swing-up and 60 recovery starts per model; 20 TIK and Standard Koopman models, 10 neural, 1 fitted linear) entering their terminal set and stabilizing within force and rail limits. Sets Ω_c are estimated using Theorem 15; formal guarantees require valid bounds. A dash indicates an exceeded budget.

	Swing-up success (%)				Recovery success (%)			
Predictor and terminal set	2 ms	1 ms	0.5 ms	0.25 ms	2 ms	1 ms	0.5 ms	0.25 ms
Fitted linear, estimated Ω_c	50	33	25	8	42	47	37	13
Neural, estimated Ω_c	91	91	91	–	97	97	97	–
Standard Koopman, latent terminal set	17	15	7	6	18	15	3	4
TIK, estimated Ω_c	100	88	88	88	100	99	99	99

6.2 Offline Stability Verification (Q2)

All 20 TIK controllers pass the estimated margin test in six regimes; none passes in the high-noise double cart-pole regime ($\epsilon = 10^{-6}$), although all have Schur physical Jacobians. No Standard Koopman controller passes. These are empirical checks, not formal certificates: the curvature estimate covers simulator curvature in only 80 of 100 near-target cart-pole fits (Table 10).

6.3 Constrained Control and Large-Amplitude Maneuvers (Q3)

Cart-pole swing-up and recovery. Table 3 tests tilts up to 135° and kicks up to 6 rad/s. Terminal sets use the Lyapunov-metric formulation in Theorem 15. All 6,982 rollouts entering Ω_c converged without subsequent violations. TIK achieves 100% swing-up success at 2 ms and 88% at 0.25 ms, where neural MPPI exceeds the budget. At 1 and 0.5 ms, neural MPPI achieves 91% and TIK achieves 88%. Tuning details are in Appendix E.1.

Double cart-pole QP-MPC. Trained on the original data only, TIK succeeds in 4% of rollouts. After adding near-target and closed-loop data it reaches 83%, compared with 0% for Standard Koopman and 16% for the same-data control without tangent matching (Appendix F.5). TIK remains below fitted linear MPC and fitted LQR (90% and 90%). Most tested starts can already be recovered by the fitted local controller.

Linear versus bilinear TIK. In the cart-pole ablation, the MPPI-only swing-up success rate is 6% with linear input and 95% with bilinear input (upright entry: 41% and 100%). With separate local balancing controllers, the bilinear model reaches 88% overall success with lifted LQR and 99% with fitted physical LQR. Their gains agree at the target, but their

actions can differ away from it (Appendix F.3). Without tangent matching, its lifted LQR policy is unstable on the plant for all 20 seeds (Table 11).

6.4 Real-Time Control on Simulated Manipulators (Q4)

The simulated seven-joint arms test control under short planning budgets (Table 12). On Kinova Gen3 obstacle avoidance (Sun et al., 2025), TIK, fitted linear MPC, and Standard Koopman each reach 100% success at 5 ms; the median TIK solve takes 2.1 ms. On the six-goal torque-controlled arm, which uses one predictor across goals with distinct linearizations, TIK reaches 100% at 1 ms, matching workspace-center linear MPC (100%) and exceeding neural NMPC (92%) and Standard Koopman (0%). This benchmark shows no success-rate advantage over linear MPC (Table 14).

7 Conclusion and Limitations

A nominal Koopman stability certificate transfers when the learned closed loop reproduces the physical closed-loop Jacobian on the embedding tangent space. TIK enforces matching to a local fit from offline transitions; the stability test bounds the remaining error between this fit and the plant. TIK recovers the fitted physical LQR gain near the target and allows state-dependent actuation in the predictor used for large maneuvers.

The formal certificates require valid curvature and noise bounds and a differentiable local model; our experiments use empirical curvature estimates in simulation. The terminal-set result covers the local controller after entry. The effects of hardware noise and contact transitions require further study.

Generative AI Use

The authors used generative AI tools to polish the writing and to help with the proofs of the theorems. The authors reviewed and verified all AI-assisted content, and they are responsible for the accuracy and integrity of the submitted work.

References

Abulikemu Abuduweili and Changliu Liu. Estimating neural network robustness via Lipschitz constant and architecture sensitivity. In *CoRL Workshop on Safe and Robust Robot Learning for Operation in the Real World*, 2024.

Abulikemu Abuduweili, Yuyang Pang, Feihan Li, and Changliu Liu. Scaling law of neural Koopman operators. *arXiv preprint arXiv:2602.19943*, 2026.

Daniel Bruder, Xun Fu, and Ram Vasudevan. Advantages of bilinear Koopman realizations for the modeling and control of systems with unknown dynamics. *IEEE Robotics and Automation Letters*, 6(3):4369–4376, 2021.

Steven L. Brunton, Bingni W. Brunton, Joshua L. Proctor, and J. Nathan Kutz. Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLoS ONE*, 11(2):e0150171, 2016.

Hongyi Chen, Abulikemu Abuduweili, Aviral Agrawal, Yunhai Han, Harish Ravichandar, Changliu Liu, and Jeffrey Ichnowski. KOROL: Learning visualizable object feature with Koopman operator rollout for manipulation. In *Conference on Robot Learning (CoRL)*, 2024.

Saber Elaydi. *An Introduction to Difference Equations*. Springer, 3 edition, 2005. doi: 10.1007/0-387-27602-5.

Fletcher Fan, Bowen Yi, David Rye, Guodong Shi, and Ian R. Manchester. Learning stable Koopman embeddings for identification and control. *Automatica*, 189:113003, 2026. doi: 10.1016/j.automatica.2026.113003.

Haizhou Ge, Mengshuo Wang, Yufei Jia, Lei Han, Lu Shi, Guyue Zhou, and Ruqi Huang. Why Koopman structure does not help differentiable world models for humanoid whole-body tracking. *arXiv preprint arXiv:2611.00951*, 2026.

Tushar Jain. Data-driven stability verification and region of attraction estimation for Koopman-lifted nonlinear systems. *IFAC Journal of Systems and Control*, 36:100424, 2026. doi: 10.1016/j.ifacsc.2026.100424.

Sebin Jung, Abulikemu Abuduweili, Jiaxing Li, and Changliu Liu. Whole-body safe control of robotic systems with Koopman neural dynamics. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 23949–23956, 2026.

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

Bernard O Koopman. Hamiltonian systems and transformation in Hilbert space. *Proceedings of the National Academy of Sciences*, 17(5):315–318, 1931.

Milan Korda and Igor Mezić. Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica*, 93:149–160, 2018.

Feihan Li, Abulikemu Abuduweili, Yifan Sun, Rui Chen, Weiye Zhao, and Changliu Liu. Continual learning and lifting of Koopman dynamics for linear control of legged robots. In *Learning for Dynamics and Control Conference (L4DC)*, volume 283 of *PMLR*, 2025.

Zexiang Liu, Necmiye Ozay, and Eduardo D. Sontag. Properties of immersions for systems with multiple limit sets with implications to learning Koopman embeddings. *Automatica*, 2025.

Bethany Lusch, J Nathan Kutz, and Steven L Brunton. Deep learning for universal linear embeddings of nonlinear dynamics. *Nature Communications*, 9(1):4950, 2018.

Giorgos Mamakoukas, Ian Abraham, and Todd D. Murphey. Learning stable models for prediction and control. *IEEE Transactions on Robotics*, 39(3):2255–2275, 2023.

Alexandre Mauroy and Igor Mezić. Global stability analysis using the eigenfunctions of the Koopman operator. *IEEE Transactions on Automatic Control*, 61(11):3356–3369, 2016. doi: 10.1109/TAC.2016.2518918.

Feliks Nüske, Sebastian Peitz, Friedrich Philipp, Manuel Schaller, and Karl Worthmann. Finite-data error bounds for Koopman-based prediction and control. *Journal of Nonlinear Science*, 33(1):14, 2023.

James M. Ortega and Werner C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. Academic Press, 1970.

Joshua L Proctor, Steven L Brunton, and J Nathan Kutz. Generalizing Koopman theory to allow for inputs and control. *SIAM Journal on Applied Dynamical Systems*, 17(1):909–930, 2018.

Li Qiu, Bo Bernhardsson, Anders Rantzer, Edward J. Davison, Peter M. Young, and John C. Doyle. A formula for computation of the real stability radius. *Automatica*, 31(6):879–890, 1995.

Xu Shang, Jorge Cortés, and Yang Zheng. On the exponential stability of Koopman model predictive control. In *Proceedings of the 8th Annual Learning for Dynamics and Control Conference*, volume 331 of *Proceedings of Machine Learning Research*, pages 1694–1712, 2026a. URL https://proceedings.mlr.press/v331/shang26a.html.

Xu Shang, Masih Haseli, Jorge Cortés, and Yang Zheng. On the existence of Koopman linear embeddings for controlled nonlinear systems. *arXiv preprint arXiv:2602.14537*, 2026b.

Haojie Shi and Max Q-H Meng. Deep Koopman operator with control for nonlinear systems. *IEEE Robotics and Automation Letters*, 7(3):7700–7707, 2022.

Robin Strässer, Manuel Schaller, Karl Worthmann, Julian Berberich, and Frank Allgöwer. Koopman-based feedback design with stability guarantees. *IEEE Transactions on Automatic Control*, 70(1):355–370, 2025.

Robin Strässer, Manuel Schaller, Karl Worthmann, Julian Berberich, and Frank Allgöwer. SafEDMD: A Koopman-based data-driven controller design framework for nonlinear dynamical systems. *Automatica*, 185:112732, 2026a. doi: 10.1016/j.automatica.2025.112732.

Robin Strässer, Karl Worthmann, Igor Mezić, Julian Berberich, Manuel Schaller, and Frank Allgöwer. An overview of Koopman-based control: From error bounds to closed-loop guarantees. *Annual Reviews in Control*, 61:101035, 2026b. doi: 10.1016/j.arcontrol.2025.101035.

Yifan Sun, Rui Chen, Kai S. Yun, Yikuan Fang, Sebin Jung, Feihan Li, Bowei Li, Weiye Zhao, and Changliu Liu. SPARK: Safe Protective and Assistive Robot Kit. In *IFAC Symposium on Robotics*, 2025.

Pieter van Goor, Robert Mahony, Manuel Schaller, and Karl Worthmann. Reprojection methods for Koopman-based modelling and prediction. In *IEEE Conference on Decision and Control (CDC)*, pages 315–321, 2023.

Kim P. Wabersich and Melanie N. Zeilinger. Linear model predictive safety certification for learning-based control. In *IEEE Conference on Decision and Control*, pages 7130–7135, 2018.

Ruigang Wang and Ian R. Manchester. Direct parameterization of Lipschitz-bounded deep networks. In *International Conference on Machine Learning (ICML)*, 2023.

Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M. Rehg, Byron Boots, and Evangelos A. Theodorou. Information theoretic MPC for model-based reinforcement learning. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, 2017.

Matthew O Williams, Ioannis G Kevrekidis, and Clarence W Rowley. A data-driven approximation of the Koopman operator: Extending dynamic mode decomposition. *Journal of Nonlinear Science*, 25(6):1307–1346, 2015.

Zhaocong Yuan, Adam W. Hall, Siqi Zhou, Lukas Brunke, Melissa Greeff, Jacopo Panerati, and Angela P. Schoellig. safe-control-gym: A unified benchmark suite for safe learning-based control and reinforcement learning in robotics. *IEEE Robotics and Automation Letters*, 7(4):11142–11149, 2022.

A Notation

Table 4: **Summary of notation.** Symbols are grouped by the section in which they are introduced. Symbols used only locally within a single proof are defined where they appear.

Symbol	Meaning
System, Koopman model, and controller (Section 3)	
$\boldsymbol{x}, \boldsymbol{u}, f$; n_x, n_u	physical state, control input, true dynamics $\boldsymbol{x}_{t+1} = f(\boldsymbol{x}_t, \boldsymbol{u}_t)$; state and input dimensions
$(\boldsymbol{x}^*, \boldsymbol{u}^*)$	target equilibrium, shifted to the origin $(\mathbf{0}, \mathbf{0})$ in the main text
$\mathbf{A}_f, \mathbf{B}_f$	Jacobians of the true dynamics f at the target equilibrium
$\mathcal{K}$	infinite-dimensional Koopman operator (Section 3)
Ψ_θ, $\Phi_\theta(\boldsymbol{x}) = [\boldsymbol{x}; \Psi_\theta(\boldsymbol{x}) - \Psi_\theta(\mathbf{0})]$, N	neural network encoder, embedding function, latent dimension
$\boldsymbol{z}$, $\hat{\boldsymbol{z}}$; $\mathbf{C} = [\mathbf{I}_{n_x}\ \mathbf{0}]$	encoded latent state, linear prediction; state projection $\boldsymbol{x} = \mathbf{C}\boldsymbol{z}$
$\mathbf{A}, \mathbf{B}$	learned Koopman matrices
$\mathbf{J}_\theta, \mathbf{J}_\theta^+$	Jacobian of Φ_θ at the target ($\mathbf{C}\mathbf{J}_\theta = \mathbf{I}$) and its Moore–Penrose pseudoinverse $(\mathbf{J}_\theta^\mathsf{T}\mathbf{J}_\theta)^{-1}\mathbf{J}_\theta^\mathsf{T}$
$\mathbf{Q}_x, \mathbf{R}$; $\mathbf{P}$; $\mathbf{K}$; N_h	LQR cost weights; Riccati terminal matrix; Koopman LQR gain (Section 3); MPC horizon
$\mathbf{A}_{\mathrm{cl}} = \mathbf{A} - \mathbf{B}\mathbf{K}$	latent closed-loop matrix; the nominal stability certificate is $\rho(\mathbf{A}_{\mathrm{cl}}) < 1$
g; $\mathbf{F} = \mathbf{A}_f - \mathbf{B}_f\mathbf{K}\mathbf{J}_\theta$	deployed nonlinear closed loop and its Jacobian at the target (3)
Tangent invariance and certificate transfer (Section 4, Appendix C.1)	
$E(\boldsymbol{x}, \boldsymbol{u})$; $\mathcal{E}$	one-step latent prediction error (Section 4); closed-loop error Jacobian (5)
$d_1(\mathbf{M})$; $r_{\mathbb{R}}(\mathbf{M})$	complex stability distance to the unit circle; real stability radius (9)
$\mathbf{\Pi} = \mathbf{J}_\theta\mathbf{C}$; $\mathbf{F}_0$, $\mathbf{\Delta}$	re-encoding projection (10); model part and physical linearization error of $\mathbf{F}$ (Proposition 9)
$\mathbf{M}$	reference matrix (Theorem 10)
TIK and data-driven certification (Section 5)	
$\boldsymbol{p}_i = [\boldsymbol{x}_i; \boldsymbol{u}_i]$; w_i; $\boldsymbol{c}, \mathbf{D}$	state-input pair of offline transition i; regression weight; offset and matrix of local fit (6)
$\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f$	offline least-squares estimate of $(\mathbf{A}_f, \mathbf{B}_f)$ from $\hat{\mathbf{D}} = [\hat{\mathbf{A}}_f\ \ \hat{\mathbf{B}}_f]$
$\mathbf{W}$; $\mathbf{N}_j$	unconstrained trainable matrix on the orthogonal complement (7); bilinear input matrices (Proposition 4(c))
$\hat{\mathbf{K}}$; $\mathbf{K}_o$; $\hat{\mathbf{F}}$	LQR gain of $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f, \mathbf{Q}_x, \mathbf{R})$; negative of the deployed policy Jacobian at the target; $\hat{\mathbf{A}}_f - \hat{\mathbf{B}}_f\mathbf{K}_o$
L, ϵ, $\mathcal{N}$	Hessian curvature constant, measurement noise bound, neighborhood in Assumption 5
ϱ_i; $\mathbf{\Sigma}$; $\boldsymbol{g}_{i,D}$; $\mathbf{\Gamma}$; $\eta(\mathbf{\Gamma})$	residual bound of transition i; weighted Gram matrix; linear estimator vector; right multiplier; error bound (8)
g_{loc}; L_H, L_T	deployed loop under local controller $\boldsymbol{u} = -\hat{\mathbf{K}}\boldsymbol{x}$; second- and third-order curvature constants (Theorem 8)
$\mathbf{S}$, $\mathbf{Q}$, $\hat{\vartheta}$, κ; r, Ω_c	Lyapunov matrix of $\hat{\mathbf{F}}$ and weight, contraction factor, condition number; certified radius and terminal set

B A Scalar Example of a Nominally Certified but Unstable Controller

Consider the scalar system $x_{t+1} = f(x_t, u_t) = 1.1x_t + 0.2x_t^2 + u_t$, which is open-loop unstable at $(x^*, u^*) = (0, 0)$ with $A_f = 1.1$ and $B_f = 1$. We use the quadratic embedding $\boldsymbol{z}_t = \Phi(x_t) = [x_t, x_t^2]^\mathsf{T}$ with Jacobian $\mathbf{J}_\theta = \partial\Phi/\partial x\big|_0 = [1, 0]^\mathsf{T}$. The true next encoded state

is $\boldsymbol{z}_{t+1} = \Phi(f(x_t, u_t)) = [1.1x_t + 0.2x_t^2 + u_t;\, (1.1x_t + 0.2x_t^2 + u_t)^2]$: because the second coordinate $z_{t+1,2} = (1.1x_t + 0.2x_t^2 + u_t)^2$ contains only terms of degree two or higher in (x_t, u_t), its first derivatives at the origin are zero. The full encoded update therefore has derivatives $\mathbf{J}_\theta A_f = [1.1,\, 0]^\mathsf{T}$ and $\mathbf{J}_\theta B_f = [1,\, 0]^\mathsf{T}$.

Consider the following approximate linear model $\hat{\boldsymbol{z}}_{t+1} = \mathbf{A}\boldsymbol{z}_t + \mathbf{B}u_t$:

$$\mathbf{A} = \begin{bmatrix} 1.1 & 0.2 \\ -0.1 & 0.9 \end{bmatrix}, \qquad \mathbf{B} = \begin{bmatrix} 1 \\ -0.75 \end{bmatrix}, \qquad \hat{\boldsymbol{z}}_{t+1} = \begin{bmatrix} 1.1x_t + 0.2x_t^2 + u_t \\ -0.1x_t + 0.9x_t^2 - 0.75u_t \end{bmatrix}.$$

Its first row $\hat{x}_{t+1} = \mathbf{C}\hat{\boldsymbol{z}}_{t+1} = 1.1x_t + 0.2x_t^2 + u_t$ reproduces $f(x_t, u_t)$ with *zero error everywhere*. The second row approximates x_{t+1}^2 using the linear terms $-0.1x_t$ and $-0.75u_t$. These terms violate tangent invariance (4) at the origin because $\mathbf{J}_\theta A_f - \mathbf{A}\mathbf{J}_\theta = [0,\, 0.1]^\mathsf{T}$ and $\mathbf{J}_\theta B_f - \mathbf{B} = [0,\, 0.75]^\mathsf{T}$.

LQR with $Q_x = R = 1$ on $(\mathbf{A}, \mathbf{B})$ gives $\mathbf{K} \approx [0.02247, -0.86339]$. In the *latent loop*, $\mathbf{A}_{\mathrm{cl}} = \mathbf{A} - \mathbf{B}\mathbf{K}$ has eigenvalues 0.95 and 0.38, so the controller is nominally certified stable with $\rho(\mathbf{A}_{\mathrm{cl}}) = 0.95 < 1$: because $\mathbf{B} = [1,\, -0.75]^\mathsf{T}$, LQR stabilizes $\mathbf{A}_{\mathrm{cl}}$ via the gain -0.863 on z_2 while placing only 0.022 on $z_1 = x_t$.

In the *deployed loop* on the true system, re-encoding $\Phi(x_t) = [x_t,\, x_t^2]^\mathsf{T}$ gives $x_{t+1} = (1.1 - K_1)x_t + (0.2 - K_2)x_t^2$. The quadratic term has zero slope at the origin, leaving $F = 1.1 - K_1 \approx 1.078 > 1$. From $x_0 = 0.01$, the latent prediction decays to 0.00274 at step 30, while the plant state grows to 1.29.

Enforcing tangent invariance sets $A_{21} = 0$ and $B_2 = 0$, giving

$$A_{\mathrm{TIK}} = \begin{bmatrix} 1.1 & 0.2 \\ 0 & 0.9 \end{bmatrix}, \qquad B_{\mathrm{TIK}} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

LQR then gives $\mathbf{K}_{\mathrm{TIK}} \approx [0.70343,\, 0.19888]$. The deployed Jacobian is $F \approx 0.397$, so the target is locally exponentially stable.

C Proofs of the Main Results

Where needed, this appendix keeps the target equilibrium $(\boldsymbol{x}^*, \boldsymbol{u}^*)$ explicit, whereas the main text shifts coordinates so that $(\boldsymbol{x}^*, \boldsymbol{u}^*) = (\mathbf{0}, \mathbf{0})$. We prove the results in Section 4 and Section 5. The proofs of Proposition 16 and Corollary 17 appear in Appendix F.4. For any real square matrix $\mathbf{M}$ with no eigenvalue on the unit circle, the real stability radius (Qiu et al., 1995) is given by

$$r_{\mathbb{R}}(\mathbf{M}) \;=\; \Big[\sup_{|s|=1} \mu_{\mathbb{R}}\big((s\mathbf{I} - \mathbf{M})^{-1}\big)\Big]^{-1}, \quad \mu_{\mathbb{R}}(\mathbf{X}) \;=\; \inf_{\gamma \in (0,1]} \sigma_2\left(\begin{bmatrix} \operatorname{Re}\mathbf{X} & -\gamma \operatorname{Im}\mathbf{X} \\ \gamma^{-1}\operatorname{Im}\mathbf{X} & \operatorname{Re}\mathbf{X} \end{bmatrix}\right), \quad (9)$$

where σ_2 denotes the second-largest singular value. A conservative numerical test uses a lower bound on the complex distance $d_1(\mathbf{M})$, which is no larger than $r_{\mathbb{R}}(\mathbf{M})$. For m equally spaced points s_j on the unit circle, the singular-value perturbation inequality gives

$$d_1(\mathbf{M}) \geq \min_j \sigma_{\min}(s_j\mathbf{I} - \mathbf{M}) - 2\sin\big(\pi/(2m)\big).$$

For a verified lower bound, subtract any numerical error bound on the sampled singular values as well. The uncorrected sampled minimum is an upper bound on d_1.

C.1 Further Results on the Deployed Loop

We first state two structural results referenced in Section 4 and Section 5. Let $\boldsymbol{r}(\boldsymbol{x}, \boldsymbol{u}) = \mathbf{C}E(\boldsymbol{x}, \boldsymbol{u})$ denote the physical state block of the one-step prediction error (Section 4), and let $\mathbf{D}_x = \partial \boldsymbol{r}/\partial \boldsymbol{x}$ and $\mathbf{D}_u = \partial \boldsymbol{r}/\partial \boldsymbol{u}$ at the target equilibrium. Define the matrix

$$\mathbf{\Pi} = \mathbf{J}_\theta \mathbf{C} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \partial \Psi_\theta / \partial \boldsymbol{x} & \mathbf{0} \end{bmatrix}. \tag{10}$$

Geometrically, $\mathbf{\Pi}$ represents state re-encoding to first order: it retains the physical state coordinates of a latent vector and replaces the encoder coordinates with the linearized re-encoding of that physical state (van Goor et al., 2023). Because $\mathbf{C}\mathbf{J}_\theta = \mathbf{I}$, $\mathbf{\Pi}$ is a projection ($\mathbf{\Pi}^2 = \mathbf{\Pi}$) whose range is the tangent space range($\mathbf{J}_\theta$).

Proposition 9 (Decomposition of the deployed Jacobian)**.** *Suppose Φ_θ is differentiable at the target equilibrium. Then:*

(a) The deployed Jacobian decomposes as $\mathbf{F} = \mathbf{F}_0 + \mathbf{\Delta}$, where $\mathbf{F}_0 = \mathbf{C}\mathbf{A}_{\text{cl}}\mathbf{J}_\theta$ is the model-predicted transition and $\mathbf{\Delta} = \mathbf{D}_x - \mathbf{D}_u\mathbf{K}\mathbf{J}_\theta$ is the physical linearization error.

(b) The matrices $\mathbf{F}_0$ and $\mathbf{\Pi}\mathbf{A}_{\text{cl}}$ share the same nonzero eigenvalues with identical multiplicities; in particular, $\rho(\mathbf{F}_0) = \rho(\mathbf{\Pi}\mathbf{A}_{\text{cl}})$.

(c) If $(\mathbf{I} - \mathbf{\Pi})\mathbf{A}_{\text{cl}}\mathbf{J}_\theta = \mathbf{0}$, meaning that $\mathbf{A}_{\text{cl}}$ maps the tangent space range($\mathbf{J}_\theta$) into itself, then the characteristic polynomial of $\mathbf{F}_0$ divides that of $\mathbf{A}_{\text{cl}}$, so $\rho(\mathbf{F}_0) \leq \rho(\mathbf{A}_{\text{cl}})$.

Proposition 9 shows that the discrepancy between $\rho(\mathbf{A}_{\text{cl}})$ and $\rho(\mathbf{F})$ arises from two distinct sources: the physical linearization error $\mathbf{\Delta}$, and the tangent-space invariance defect $(\mathbf{I} - \mathbf{\Pi})\mathbf{A}_{\text{cl}}\mathbf{J}_\theta$ when the latent loop leaves the tangent space. Even a model with zero physical linearization error $\mathbf{\Delta} = \mathbf{0}$ can fail on the true system if $(\mathbf{I} - \mathbf{\Pi})\mathbf{A}_{\text{cl}}\mathbf{J}_\theta \neq \mathbf{0}$, which can make $\rho(\mathbf{\Pi}\mathbf{A}_{\text{cl}}) > 1$ while $\rho(\mathbf{A}_{\text{cl}}) < 1$.

Theorem 10 (Stability from a reference matrix)**.** *Assume a differentiable lifting and inactive input limits near the target. Let $\mathbf{M}$ be a real $n_x \times n_x$ matrix with no eigenvalue on the unit circle. If $\|\mathbf{F} - \mathbf{M}\|_2 < r_{\mathbb{R}}(\mathbf{M})$, then $\mathbf{F}$ and $\mathbf{M}$ have the same number of eigenvalues outside the closed unit disc and none on the unit circle. In particular:*

(a) Taking $\mathbf{M} = \mathbf{F}_0$, if $\|\mathbf{\Delta}\|_2 < r_{\mathbb{R}}(\mathbf{F}_0)$, local exponential stability of the deployed loop is decided by $\rho(\mathbf{F}_0) = \rho(\mathbf{\Pi}\mathbf{A}_{\text{cl}}) < 1$.

(b) Taking $\mathbf{M} = \hat{\mathbf{F}} = \hat{\mathbf{A}}_f - \hat{\mathbf{B}}_f\mathbf{K}_o$ constructed from estimated Jacobians for any feedback policy preserving the target equilibrium and differentiable there with Jacobian $-\mathbf{K}_o$, where $\mathbf{K}_o = \mathbf{K}\mathbf{J}_\theta$ for lifted LQR and $\mathbf{K}_o = \hat{\mathbf{K}}$ for TIK, the condition holds whenever $\|(\hat{\mathbf{D}} - [\mathbf{A}_f \;\; \mathbf{B}_f])[\mathbf{I}; -\mathbf{K}_o]\|_2 < r_{\mathbb{R}}(\hat{\mathbf{F}})$.

C.2 Proofs of Transfer, Parameterization, and Stability

Proof of Proposition 9. (a) Differentiating $f(\boldsymbol{x}, \boldsymbol{u}) = \mathbf{C}(\mathbf{A}\Phi_\theta(\boldsymbol{x}) + \mathbf{B}\boldsymbol{u}) + \boldsymbol{r}(\boldsymbol{x}, \boldsymbol{u})$ at the target yields $\mathbf{A}_f = \mathbf{C}\mathbf{A}\mathbf{J}_\theta + \mathbf{D}_x$ and $\mathbf{B}_f = \mathbf{C}\mathbf{B} + \mathbf{D}_u$. Substituting these into $\mathbf{F} = \mathbf{A}_f - \mathbf{B}_f\mathbf{K}\mathbf{J}_\theta$ gives $\mathbf{F} = \mathbf{C}(\mathbf{A} - \mathbf{B}\mathbf{K})\mathbf{J}_\theta + (\mathbf{D}_x - \mathbf{D}_u\mathbf{K}\mathbf{J}_\theta) = \mathbf{F}_0 + \mathbf{\Delta}$. (b) Because $\mathbf{C}\mathbf{J}_\theta = \mathbf{I}$, we have $\mathbf{\Pi}^2 = \mathbf{J}_\theta(\mathbf{C}\mathbf{J}_\theta)\mathbf{C} = \mathbf{\Pi}$ and range($\mathbf{\Pi}$) = range($\mathbf{J}_\theta$). For any matrices $\mathbf{X} \in \mathbb{R}^{N \times n_x}$ and $\mathbf{Y} \in \mathbb{R}^{n_x \times N}$, the products $\mathbf{X}\mathbf{Y}$ and $\mathbf{Y}\mathbf{X}$ share the same nonzero eigenvalues with identical algebraic multiplicities; setting $\mathbf{X} = \mathbf{J}_\theta$ and $\mathbf{Y} = \mathbf{C}\mathbf{A}_{\text{cl}}$ gives $\mathbf{Y}\mathbf{X} = \mathbf{F}_0$ and $\mathbf{X}\mathbf{Y} = \mathbf{\Pi}\mathbf{A}_{\text{cl}}$. (c) The condition $(\mathbf{I} - \mathbf{\Pi})\mathbf{A}_{\text{cl}}\mathbf{J}_\theta = \mathbf{0}$ implies $\mathbf{A}_{\text{cl}}\mathbf{J}_\theta = \mathbf{\Pi}\mathbf{A}_{\text{cl}}\mathbf{J}_\theta = \mathbf{J}_\theta\mathbf{F}_0$. Because $\mathbf{J}_\theta$ has full column rank, range($\mathbf{J}_\theta$) is an $\mathbf{A}_{\text{cl}}$-invariant subspace and $\mathbf{F}_0$ is the matrix representation of

$\mathbf{A}_{\mathrm{cl}}$ restricted to range($\mathbf{J}_\theta$) in the basis formed by the columns of $\mathbf{J}_\theta$. Extending the columns of $\mathbf{J}_\theta$ to a full basis of $\mathbb{R}^N$ puts $\mathbf{A}_{\mathrm{cl}}$ into block upper-triangular form with $\mathbf{F}_0$ as its top-left diagonal block; hence the characteristic polynomial of $\mathbf{F}_0$ divides that of $\mathbf{A}_{\mathrm{cl}}$, and every eigenvalue of $\mathbf{F}_0$ is an eigenvalue of $\mathbf{A}_{\mathrm{cl}}$. □

Proof of Theorem 2. For $\mathbf{F}\boldsymbol{v} = \lambda\boldsymbol{v}$ with $\|\boldsymbol{v}\|_2 = 1$, (5) gives $(\lambda\mathbf{I} - \mathbf{A}_{\mathrm{cl}})\mathbf{J}_\theta\boldsymbol{v} = \mathcal{E}\boldsymbol{v}$. Since $\|\mathbf{J}_\theta\boldsymbol{v}\|_2 \geq \|\mathbf{C}\mathbf{J}_\theta\boldsymbol{v}\|_2 = 1$, $\sigma_{\min}(\lambda\mathbf{I} - \mathbf{A}_{\mathrm{cl}}) \leq \|\mathcal{E}\boldsymbol{v}\|_2/\|\mathbf{J}_\theta\boldsymbol{v}\|_2 \leq \|\mathcal{E}\|_2$. If $\mathcal{E} = \mathbf{0}$, $\mathbf{J}_\theta\boldsymbol{v}$ is a nonzero eigenvector of $\mathbf{A}_{\mathrm{cl}}$ with eigenvalue λ. For part (b), when $\mathbf{A}_{\mathrm{cl}}$ is Schur, the resolvent norm $s \mapsto \|(s\mathbf{I} - \mathbf{A}_{\mathrm{cl}})^{-1}\|_2$ is subharmonic on the exterior of the open unit disc $\{s \in \mathbb{C} : |s| \geq 1\}$ and vanishes as $|s| \to \infty$. By the maximum modulus principle for subharmonic functions, it attains its maximum on the boundary $|s| = 1$. Taking reciprocals gives $\min_{|s|\geq 1} \sigma_{\min}(s\mathbf{I} - \mathbf{A}_{\mathrm{cl}}) = \min_{|s|=1} \sigma_{\min}(s\mathbf{I} - \mathbf{A}_{\mathrm{cl}}) = d_1(\mathbf{A}_{\mathrm{cl}})$. Thus, if $\|\mathcal{E}\|_2 < d_1(\mathbf{A}_{\mathrm{cl}})$, no eigenvalue of $\mathbf{F}$ can satisfy $|\lambda| \geq 1$, so $\rho(\mathbf{F}) < 1$ and Lyapunov's indirect method (Khalil, 2002) ensures local exponential stability of the deployed loop. □

Proof of Corollary 3. If $\mathbf{A}_f\boldsymbol{v} = \lambda\boldsymbol{v}$ with $\boldsymbol{v} \neq \mathbf{0}$, then $\mathbf{A}\mathbf{J}_\theta\boldsymbol{v} = \lambda\mathbf{J}_\theta\boldsymbol{v}$ by (4). Full column rank gives $\mathbf{J}_\theta\boldsymbol{v} \neq \mathbf{0}$, so $\rho(\mathbf{A}) \geq \rho(\mathbf{A}_f) > 1$. □

Proof of Theorem 10. Let $\mathbf{Z} = \mathbf{F} - \mathbf{M} \in \mathbb{R}^{n_x \times n_x}$. Consider the continuous matrix path $\mathbf{M}(\tau) = \mathbf{M} + \tau\mathbf{Z}$ for $\tau \in [0, 1]$. If $\mathbf{M}(\tau)$ had an eigenvalue s on the unit circle ($|s| = 1$) for some $\tau \in (0, 1]$, then $s\mathbf{I} - \mathbf{M} - \tau\mathbf{Z} = (s\mathbf{I} - \mathbf{M})(\mathbf{I} - \tau(s\mathbf{I} - \mathbf{M})^{-1}\mathbf{Z})$ would be singular. By the definition of the real structured singular value $\mu_{\mathbb{R}}$ (Qiu et al., 1995), this requires $\tau\|\mathbf{Z}\|_2 \geq 1/\mu_{\mathbb{R}}((s\mathbf{I} - \mathbf{M})^{-1}) \geq r_{\mathbb{R}}(\mathbf{M})$, which contradicts $\|\mathbf{Z}\|_2 < r_{\mathbb{R}}(\mathbf{M})$. Because the eigenvalues of $\mathbf{M}(\tau)$ vary continuously with τ and never cross the unit circle, $\mathbf{F} = \mathbf{M}(1)$ and $\mathbf{M} = \mathbf{M}(0)$ have the same number of eigenvalues outside the unit disc. Part (a) follows with $\mathbf{M} = \mathbf{F}_0$ and Proposition 9(b); part (b) follows from $\mathbf{F} - \hat{\mathbf{F}} = ([\mathbf{A}_f \;\; \mathbf{B}_f] - \hat{\mathbf{D}})[\mathbf{I}; -\mathbf{K}_o]$. □

Proof of Proposition 4. We prove the result in the general coordinate form $\mathbf{G} = \mathbf{J}_\theta\mathbf{T}$, where $\mathbf{T} = \mathbf{I}$ in standard coordinates and $\mathbf{T}$ is the chart Jacobian when an angle θ is represented in the state by $(\sin\theta, \cos\theta)$ (in which case $\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f$ are expressed in minimal angle coordinates and $\mathbf{Q}_x$ is replaced by $\mathbf{T}^\mathsf{T}\mathbf{Q}_x\mathbf{T}$). In this coordinate form, replace $\mathbf{J}_\theta$ by $\mathbf{G}$ in (7).

(a) Because $\mathbf{G}$ has full column rank, $\mathbf{G}^+\mathbf{G} = \mathbf{I}$ and $(\mathbf{I} - \mathbf{G}\mathbf{G}^+)\mathbf{G} = \mathbf{0}$. Multiplying (7) on the right by $\mathbf{G}$ yields $\mathbf{A}\mathbf{G} = \mathbf{G}\hat{\mathbf{A}}_f$, while $\mathbf{B} = \mathbf{G}\hat{\mathbf{B}}_f$ holds by definition. Substituting these into (5) yields $\mathcal{E} = \mathbf{G}\mathbf{F} - (\mathbf{A} - \mathbf{B}\mathbf{K})\mathbf{G} = \mathbf{G}(\mathbf{A}_f - \hat{\mathbf{A}}_f) - \mathbf{G}(\mathbf{B}_f - \hat{\mathbf{B}}_f)\mathbf{K}\mathbf{G}$.

(b) In a basis of $\mathbb{R}^N$ adapted to range($\mathbf{G}$) $\oplus$ range($\mathbf{G}$)$^\perp$, $\mathbf{A}$ is block upper-triangular with top-left block $\hat{\mathbf{A}}_f$ and bottom-right block given by the restriction of $(\mathbf{I} - \mathbf{G}\mathbf{G}^+)\mathbf{W}(\mathbf{I} - \mathbf{G}\mathbf{G}^+)$ to range($\mathbf{G}$)$^\perp$, while range($\mathbf{B}$) $\subseteq$ range($\mathbf{G}$) lies entirely in the first block. Because the orthogonal complement block is Schur by assumption and $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ is stabilizable, the full latent pair $(\mathbf{A}, \mathbf{B})$ is stabilizable. Similarly, any eigenvector of $\mathbf{A}$ with $|\lambda| \geq 1$ must lie in range($\mathbf{G}$), where $\mathbf{C}\mathbf{G} = \mathbf{T}$ has full column rank, so detectability of $(\hat{\mathbf{A}}_f, (\mathbf{T}^\mathsf{T}\mathbf{Q}_x\mathbf{T})^{1/2})$ ensures detectability of $(\mathbf{A}, \mathbf{Q}_x^{1/2}\mathbf{C})$. Therefore, the latent LQR Riccati equation has a unique stabilizing solution $\mathbf{P} \succeq 0$. For an initial state on the tangent space, $\boldsymbol{z}_0 = \mathbf{G}\boldsymbol{x}_0$, any control sequence $\boldsymbol{u}_{0:\infty}$ keeps the latent trajectory inside range($\mathbf{G}$) via $\boldsymbol{z}_t = \mathbf{G}\boldsymbol{x}_t$ with $\boldsymbol{x}_{t+1} = \hat{\mathbf{A}}_f\boldsymbol{x}_t + \hat{\mathbf{B}}_f\boldsymbol{u}_t$, and the stage costs match identically: $\boldsymbol{z}_t^\mathsf{T}\mathbf{C}^\mathsf{T}\mathbf{Q}_x\mathbf{C}\boldsymbol{z}_t = \boldsymbol{x}_t^\mathsf{T}\mathbf{T}^\mathsf{T}\mathbf{Q}_x\mathbf{T}\boldsymbol{x}_t$. Because $\mathbf{R} \succ 0$ makes the optimal control sequence unique, the optimal first control action of the latent LQR problem from $\boldsymbol{z}_0 = \mathbf{G}\boldsymbol{x}_0$ must equal that of the physical LQR problem on $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f, \mathbf{T}^\mathsf{T}\mathbf{Q}_x\mathbf{T}, \mathbf{R})$: $-\mathbf{K}\mathbf{G}\boldsymbol{x}_0 = -\hat{\mathbf{K}}\boldsymbol{x}_0$ for all $\boldsymbol{x}_0$, proving $\mathbf{K}\mathbf{G} = \hat{\mathbf{K}}$.

(c) With the target equilibrium written explicitly, the bilinear term is $\sum_{j=1}^{n_u}(u_{t,j} - u_j^*)\mathbf{N}_j(\boldsymbol{z}_t - \Phi_\theta(\boldsymbol{x}^*))$. Both this term and its first partial derivatives with respect to $\boldsymbol{z}_t$ and $\boldsymbol{u}_t$ vanish at $(\boldsymbol{z}_t, \boldsymbol{u}_t) = (\Phi_\theta(\boldsymbol{x}^*), \boldsymbol{u}^*)$, leaving the linearization $(\mathbf{A}, \mathbf{B})$ and properties (a)–(b) unchanged. □

Proof of Proposition 7. Since $\mathbf{R} \succ 0$, substituting $\boldsymbol{z}_k = \mathbf{A}^k \boldsymbol{z}_0 + \sum_{j=0}^{k-1} \mathbf{A}^{k-1-j} \mathbf{B} \boldsymbol{u}_j$ turns (2) into a strictly convex quadratic program in the input sequence $\boldsymbol{u}_{0:N_h-1}$. Since the target equilibrium lies in the strict interior of the state and input constraints, the unconstrained minimizer satisfies all constraints in an open neighborhood of the target. With the stabilizing Riccati solution $\mathbf{P}$ as the terminal weight, the terminal cost $\boldsymbol{z}_{N_h}^\mathsf{T} \mathbf{P} \boldsymbol{z}_{N_h}$ equals the optimal infinite-horizon LQR cost-to-go from $\boldsymbol{z}_{N_h}$, so by Bellman's principle of optimality the unconstrained finite-horizon minimizer applies $\boldsymbol{u}_0 = -\mathbf{K}\boldsymbol{z}_0$ at the first step for every horizon $N_h \geq 1$. Combining this with Proposition 4(b) completes the proof. □

Proof of Theorem 6. In the centered coordinates of the main text, write $\mathbf{H}_f = [\mathbf{D}_f \ \mathbf{0}]$ and $\hat{\mathbf{H}} = [\hat{\mathbf{D}} \ \hat{\boldsymbol{c}}]$. Since $\boldsymbol{x}_{i+1} = \mathbf{H}_f \boldsymbol{q}_i + \boldsymbol{\ell}_i$, the weighted normal equations give

$$(\hat{\mathbf{H}} - \mathbf{H}_f)\boldsymbol{\Sigma} = \sum_i w_i \boldsymbol{\ell}_i \boldsymbol{q}_i^\mathsf{T}, \qquad (\hat{\mathbf{D}} - \mathbf{D}_f)\boldsymbol{\Gamma} = \sum_i w_i \boldsymbol{\ell}_i (\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D})^\mathsf{T}.$$

The full-rank condition makes it possible to invert $\boldsymbol{\Sigma}$. Bounding each rank-one summand by $w_i \varrho_i \|\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D}\|_2$ proves (8), including the fitted intercept. This deterministic inequality holds simultaneously for all dimensionally compatible $\boldsymbol{\Gamma}$, so a gain chosen using the same data is allowed when the curvature and noise bounds hold. The deployed map $g(\boldsymbol{x}) = f(\boldsymbol{x}, \boldsymbol{\pi}(\boldsymbol{x}))$ is differentiable at the equilibrium with Jacobian $\mathbf{F} = \mathbf{A}_f - \mathbf{B}_f \mathbf{K}_o = [\mathbf{A}_f \ \ \mathbf{B}_f]\boldsymbol{\Gamma}$, while $\hat{\mathbf{F}} = \hat{\mathbf{D}}\boldsymbol{\Gamma}$ with $\boldsymbol{\Gamma} = [\mathbf{I}; -\mathbf{K}_o]$. By (8), $\|\mathbf{F} - \hat{\mathbf{F}}\|_2 \leq \eta < r_\mathbb{R}(\hat{\mathbf{F}})$. Applying Theorem 10(b) and Lyapunov's indirect method (Khalil, 2002) establishes local exponential stability when $\rho(\hat{\mathbf{F}}) < 1$ and instability when $\rho(\hat{\mathbf{F}}) > 1$. □

The following two corollaries extend Theorem 6 to stochastic measurement noise (including noisy regression states) and to third-order Taylor remainders used in our two-scale offline verification protocol (Section 6.2).

Corollary 11 (Offline certificate with stochastic and regressor noise)**.** *In Theorem 6, condition on the regressors, weights, and* $\boldsymbol{\Gamma}$*. Suppose the additional successor errors are conditionally independent across transitions, conditionally zero mean, and bounded by* $|e_{ij}| \leq \epsilon'$*. This conditioning must preserve the noise assumptions; a gain fitted to the same noisy successors does not automatically satisfy them. Then with probability at least* $1 - \delta$*, the conclusion of Theorem 6 holds with* η *replaced by*

$$\eta_\delta \ = \ \eta \ + \ \epsilon' \sqrt{2 n_x \log(2 n_x m/\delta)} \left(\sum_i w_i^2 \|\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D}\|_2^2 \right)^{1/2}, \tag{11}$$

Here m *is the number of columns of* $\boldsymbol{\Gamma}$*. For bounded regressor errors* $\tilde{\boldsymbol{p}}_i = \boldsymbol{p}_i + \boldsymbol{v}_i$*,* $\|\boldsymbol{v}_i\|_2 \leq b$*, use the full-rank measured design to compute the Gram matrix and estimator vectors. In* ϱ_i*, replace* ϵ *by* $\epsilon + D_{\max} b$ *for a valid bound* $D_{\max} \geq \|Df(\boldsymbol{p}^*)\|_2$*, and upper-bound the distance by* $\|\tilde{\boldsymbol{p}}_i - \boldsymbol{p}^*\|_2 + b$*. The deterministic regressor correction also applies without the stochastic term.*

Proof. The zero-mean noise contribution to $(\hat{\mathbf{D}} - \mathbf{D}_f)\boldsymbol{\Gamma}$ is $\boldsymbol{\Xi} = \sum_i w_i \boldsymbol{e}_i (\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D})^\mathsf{T}$. Each entry Ξ_{jl} is a sum of independent zero-mean variables bounded by $\epsilon' w_i |(\boldsymbol{\Gamma}^\mathsf{T} \boldsymbol{g}_{i,D})_l|$. Applying Hoeffding's inequality to each of the $n_x m$ entries with failure probability $\delta/(n_x m)$ and bounding $\|\boldsymbol{\Xi}\|_2 \leq \|\boldsymbol{\Xi}\|_F$ yields (11). The regressor correction follows by expanding at the true target and writing the linear term as $Df(\boldsymbol{p}^*)(\tilde{\boldsymbol{p}}_i - \boldsymbol{p}^*) - Df(\boldsymbol{p}^*)\boldsymbol{v}_i$. For a gain fitted to the same noisy successors, the conditional independence in the corollary needs separate justification. One can instead bound the estimation error with $\boldsymbol{\Gamma} = \mathbf{I}$ and multiply by $\|\boldsymbol{\Gamma}\|_2$ for the fitted gain. □

Corollary 12 (Offline certificate with a cubic Taylor remainder)**.** *Let f be three times continuously differentiable (C^3) on the neighborhood $\mathcal{N}$ in Assumption 5. In the induced Euclidean multilinear norms, assume $\|D^2 f(\boldsymbol{p}^*)\| \le L_H$ and $\|D^3 f(\boldsymbol{p})\| \le L_T$ for all $\boldsymbol{p} \in \mathcal{N}$. Then the estimation error bound (8) and Theorem 6 hold with $\varrho_i = \frac{1}{2} L_H \|\boldsymbol{p}_i - \boldsymbol{p}^*\|_2^2 + \frac{1}{6} L_T \|\boldsymbol{p}_i - \boldsymbol{p}^*\|_2^3 + \epsilon$. The same argument applies in an input-weighted norm when the derivative bounds and Taylor distances use that norm.*

Proof. By Taylor's theorem with integral remainder around $\boldsymbol{p}^*$, for $\boldsymbol{d} = \boldsymbol{p}_i - \boldsymbol{p}^*$ we have $f(\boldsymbol{p}_i) - f(\boldsymbol{p}^*) - Df(\boldsymbol{p}^*)\boldsymbol{d} = \frac{1}{2} D^2 f(\boldsymbol{p}^*)[\boldsymbol{d}, \boldsymbol{d}] + \int_0^1 \frac{(1-t)^2}{2} D^3 f(\boldsymbol{p}^* + t\boldsymbol{d})[\boldsymbol{d}, \boldsymbol{d}, \boldsymbol{d}]\, dt$, whose Euclidean norm is bounded by $\frac{1}{2} L_H \|\boldsymbol{d}\|_2^2 + \frac{1}{6} L_T \|\boldsymbol{d}\|_2^3$. Adding the measurement noise bound ϵ gives the stated residual bound ϱ_i. □

C.3 Remarks on the Main Results

Why first-order tangent invariance. Vanishing first derivatives of the one-step latent error E are sufficient for exact transfer in Theorem 2. Without first-order tangent invariance, a Schur $\mathbf{A}_{\mathrm{cl}}$ places no bound on $\rho(\mathbf{F}_0)$: partitioning $\mathbf{A}_{\mathrm{cl}} = \left[\begin{smallmatrix} \mathbf{A}_{xx} & \mathbf{A}_{x\Psi} \\ \mathbf{A}_{\Psi x} & \mathbf{A}_{\Psi\Psi} \end{smallmatrix}\right]$ gives $\mathbf{F}_0 = \mathbf{A}_{xx} + \mathbf{A}_{x\Psi}\mathbf{J}_{\Psi_\theta}$. Whenever $\mathbf{A}_{x\Psi} \neq \mathbf{0}$, choosing $\mathbf{J}_{\Psi_\theta} = \omega\, \boldsymbol{v}_1 \boldsymbol{u}_1^\top$ along the leading singular vectors of $\mathbf{A}_{x\Psi}$ makes $\rho(\mathbf{F}_0) \ge |\mathrm{tr}\mathbf{F}_0|/n_x = |\mathrm{tr}\mathbf{A}_{xx} + \omega\sigma_1(\mathbf{A}_{x\Psi})|/n_x$ grow arbitrarily large with ω while $\mathbf{A}_{\mathrm{cl}}$ remains fixed and Schur.

Matching for one gain or all gains. Theorem 2(a) requires the closed-loop error Jacobian $\mathcal{E} = \mathbf{J}_\theta \mathbf{F} - \mathbf{A}_{\mathrm{cl}} \mathbf{J}_\theta$ to vanish for the specific feedback gain $\mathbf{K}$ in use. Tangent invariance (Definition 1) is an open-loop property of the learned model $(\mathbf{A}, \mathbf{B})$ alone ($\mathbf{J}_\theta \mathbf{A}_f = \mathbf{A}\mathbf{J}_\theta$ and $\mathbf{J}_\theta \mathbf{B}_f = \mathbf{B}$), which by (5) guarantees $\mathcal{E} = \mathbf{0}$ simultaneously for every feedback gain $\mathbf{K}$. Given offline estimates $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$, one can also impose tangent invariance after training via the minimum-norm projection

$$\mathbf{A}' \;=\; \mathbf{A} + (\mathbf{J}_\theta \hat{\mathbf{A}}_f - \mathbf{A}\mathbf{J}_\theta)\mathbf{J}_\theta^+, \qquad \mathbf{B}' \;=\; \mathbf{J}_\theta \hat{\mathbf{B}}_f, \tag{12}$$

which we term *anchored invariance*. However, modifying $(\mathbf{A}, \mathbf{B})$ after training changes the LQR gain $\mathbf{K}$; for uncentered controllers this shifts the closed-loop equilibrium, requiring iterative re-anchoring or locating the equilibrium from offline data via the Newton–Kantorovich theorem (Ortega and Rheinboldt, 1970). TIK in Proposition 4 avoids this iteration by centering at the target equilibrium and enforcing tangent invariance by construction throughout training.

Estimating curvature constants and selecting LQR weights. In our two-scale offline verification protocol of Section 6.2, fitting a local cubic polynomial on a wider neighborhood averages out observation noise; at upright pendulum equilibria, odd symmetry implies $D^2 f(\mathbf{0}, \mathbf{0}) = \mathbf{0}$, making a cubic fit necessary to capture the leading nonlinear remainder. The empirical constants L_H and L_T are computed from the scaled norms of the fitted second- and third-order polynomial tensors plus their Hoeffding noise bounds, multiplied by a safety factor of 2. Because $\hat{\mathbf{K}} = \mathbf{K}\mathbf{J}_\theta$ is the LQR gain of $(\hat{\mathbf{A}}_f, \hat{\mathbf{B}}_f)$ for any choice of cost weights $(\mathbf{Q}_x, \mathbf{R})$ without retraining the TIK model, the weights can be chosen to enlarge the real stability radius $r_{\mathbb{R}}(\hat{\mathbf{F}})$ relative to the estimation bound η, as detailed in Appendix F.2.

D Proof of the Certified Terminal Set

We first prove Theorem 15, which measures distances and curvature in the Lyapunov metric defined by $\mathbf{S} \succ 0$. We then obtain the Euclidean bound in Theorem 8. The constrained

swing-up experiments in Section 6.3 use the metric formulation.

Let the true nonlinear system be $\boldsymbol{x}_{t+1} = f(\boldsymbol{x}_t, \boldsymbol{u}_t)$ with target equilibrium $f(\boldsymbol{x}^*, \boldsymbol{u}^*) = \boldsymbol{x}^*$, input box $\mathcal{U} = \{\boldsymbol{u} : \boldsymbol{u}_{\min} \le \boldsymbol{u} \le \boldsymbol{u}_{\max}\}$ containing $\boldsymbol{u}^*$ in its interior, and state box $\mathcal{X} = \{\boldsymbol{x} : \boldsymbol{x}_{\mathrm{lo}} \le \boldsymbol{x} \le \boldsymbol{x}_{\mathrm{hi}}\}$ containing $\boldsymbol{x}^*$ in its interior. Write $\bar{u}_i = \min(u_{\max,i} - u_i^*,\, u_i^* - u_{\min,i}) > 0$ and $\boldsymbol{e} = \boldsymbol{x} - \boldsymbol{x}^*$. Under the local linear feedback law

$$\boldsymbol{\pi}(\boldsymbol{x}) = \boldsymbol{u}^* - \hat{\mathbf{K}}(\boldsymbol{x} - \boldsymbol{x}^*), \tag{13}$$

the input satisfies its limits on the set $\mathcal{S} = \{\boldsymbol{x} : |\hat{\mathbf{K}}_i(\boldsymbol{x} - \boldsymbol{x}^*)| \le \bar{u}_i \ \forall i\}$, where the deployed closed-loop map is $g_{\mathrm{loc}}(\boldsymbol{x}) = f(\boldsymbol{x}, \boldsymbol{\pi}(\boldsymbol{x}))$ with Jacobian $\mathbf{F} = Dg_{\mathrm{loc}}(\boldsymbol{x}^*) = \mathbf{A}_f - \mathbf{B}_f\hat{\mathbf{K}}$. Given $\mathbf{S} \succ 0$, let $\mathbf{V} = \mathbf{S}^{1/2}$, define the weighted vector norm $\|\boldsymbol{a}\|_{\mathbf{S}} = \|\mathbf{V}\boldsymbol{a}\|_2$, and denote the induced matrix norm by $\|\mathbf{M}\|_{\mathbf{S}} = \|\mathbf{V}\mathbf{M}\mathbf{V}^{-1}\|_2$. All terms in the metric test must use these transformed coordinates: the error bound controls $\mathbf{V}(\hat{\mathbf{F}} - \mathbf{F})\mathbf{V}^{-1}$, and a derivative of order k transforms as $\mathbf{V}D^k g_{\mathrm{loc}}[\mathbf{V}^{-1}\cdot, \ldots, \mathbf{V}^{-1}\cdot]$. Euclidean curvature estimates cannot be substituted without the corresponding norm factors.

Assumption 13 (Closed-loop curvature in the Lyapunov metric)**.** Let g_{loc} be C^3 on a convex set $\mathcal{N} \subseteq \mathcal{S}$ containing the target and all regressors. In the multilinear norms induced by $\|\cdot\|_{\mathbf{S}}$, assume $\|D^2 g_{\mathrm{loc}}(\boldsymbol{x}^*)\|_{\mathbf{S}} \le L_H^{\mathbf{S}}$ and $\|D^3 g_{\mathrm{loc}}(\boldsymbol{x})\|_{\mathbf{S}} \le L_T^{\mathbf{S}}$ for all $\boldsymbol{x} \in \mathcal{N}$. Offline closed-loop transitions satisfy $\boldsymbol{y}_i = g_{\mathrm{loc}}(\boldsymbol{x}_i) + \boldsymbol{\epsilon}_i$ with $\|\boldsymbol{\epsilon}_i\|_2 \le \epsilon$.

Lemma 14 (Metric Taylor remainder and least-squares fit bound)**.** *Under Assumption 13:*

(a) For every $\boldsymbol{x} \in \mathcal{N}$ with $\boldsymbol{e} = \boldsymbol{x} - \boldsymbol{x}^*$, $\|g_{\mathrm{loc}}(\boldsymbol{x}) - \boldsymbol{x}^* - \mathbf{F}\boldsymbol{e}\|_{\mathbf{S}} \le \frac{1}{2}L_H^{\mathbf{S}}\|\boldsymbol{e}\|_{\mathbf{S}}^2 + \frac{1}{6}L_T^{\mathbf{S}}\|\boldsymbol{e}\|_{\mathbf{S}}^3$.

(b) If $\hat{\mathbf{F}}$ is fitted by least squares on pairs $(\boldsymbol{x}_i, \boldsymbol{y}_i)$ in $\mathcal{N}$ with regressors $\boldsymbol{\zeta}_i = [\boldsymbol{e}_i; 1]$, Gram matrix $\boldsymbol{\Sigma} = \sum_i \boldsymbol{\zeta}_i\boldsymbol{\zeta}_i^{\mathsf{T}} \succ 0$, and $\boldsymbol{g}_{i,D}$ denoting the first n_x entries of $\boldsymbol{\Sigma}^{-1}\boldsymbol{\zeta}_i$, then

$$\|\hat{\mathbf{F}} - \mathbf{F}\|_{\mathbf{S}} \ \le\ \eta_{\mathbf{S}} \ :=\ \min\Big\{\sum_i \varrho_i\|\mathbf{V}^{-1}\boldsymbol{g}_{i,D}\|_2,\ \sqrt{n_x}\max_j \sum_i \varrho_i|(\mathbf{V}^{-1}\boldsymbol{g}_{i,D})_j|\Big\}, \tag{14}$$

where $\varrho_i = \frac{1}{2}L_H^{\mathbf{S}}\|\boldsymbol{e}_i\|_{\mathbf{S}}^2 + \frac{1}{6}L_T^{\mathbf{S}}\|\boldsymbol{e}_i\|_{\mathbf{S}}^3 + \|\mathbf{V}\|_2\,\epsilon$.

Proof. Part (a) follows from Taylor's theorem with integral remainder along the segment from $\boldsymbol{x}^*$ to $\boldsymbol{x}$ in the norm $\|\cdot\|_{\mathbf{S}}$. For part (b), writing $\boldsymbol{y}_i = \boldsymbol{x}^* + \mathbf{F}\boldsymbol{e}_i + \boldsymbol{\ell}_i$ gives $\mathbf{V}(\hat{\mathbf{F}} - \mathbf{F})\mathbf{V}^{-1} = \sum_i(\mathbf{V}\boldsymbol{\ell}_i)(\mathbf{V}^{-1}\boldsymbol{g}_{i,D})^{\mathsf{T}}$ with $\|\mathbf{V}\boldsymbol{\ell}_i\|_2 \le \varrho_i$ by part (a). For any unit vector $\boldsymbol{b} \in \mathbb{R}^{n_x}$, the triangle inequality gives $\big\|\sum_i(\mathbf{V}\boldsymbol{\ell}_i)(\mathbf{V}^{-1}\boldsymbol{g}_{i,D})^{\mathsf{T}}\boldsymbol{b}\big\|_2 \le \sum_i \varrho_i|(\mathbf{V}^{-1}\boldsymbol{g}_{i,D})^{\mathsf{T}}\boldsymbol{b}|$; applying Cauchy–Schwarz yields the first bound in (14), and applying $\|\boldsymbol{b}\|_1 \le \sqrt{n_x}$ yields the second. The weighted Euclidean counterpart is derived in the proof of Theorem 6. □

Theorem 15 (Certified invariant terminal set in the Lyapunov metric)**.** *Let $\mathbf{S} \succ 0$, suppose Assumption 13 holds, and let $\hat{\mathbf{F}}$ and $\eta_{\mathbf{S}}$ be given by Lemma 14. Set $\hat{\vartheta}_{\mathbf{S}} = \|\hat{\mathbf{F}}\|_{\mathbf{S}}$ and assume $\delta := 1 - \hat{\vartheta}_{\mathbf{S}} - \eta_{\mathbf{S}} > 0$. Let s^* be the positive root of $\frac{1}{2}L_H^{\mathbf{S}}s + \frac{1}{6}L_T^{\mathbf{S}}s^2 = \delta$, or $+\infty$ if both curvature bounds are zero. Define*

$$c^* = \min\Big\{(s^*)^2,\ c_{\mathcal{N}},\ \min_i \frac{\bar{u}_i^2}{\hat{\mathbf{K}}_i\mathbf{S}^{-1}\hat{\mathbf{K}}_i^{\mathsf{T}}},\ \min_{j,\pm}\frac{(b_j^{\pm})^2}{(\mathbf{S}^{-1})_{jj}}\Big\}, \tag{15}$$

where $c_{\mathcal{N}} = \sup\{c : \Omega_c \subseteq \mathcal{N}\}$, $\Omega_c = \{\boldsymbol{x} : (\boldsymbol{x} - \boldsymbol{x}^*)^{\mathsf{T}}\mathbf{S}(\boldsymbol{x} - \boldsymbol{x}^*) \le c\}$, $b_j^+ = x_{\mathrm{hi},j} - x_j^*$, and $b_j^- = x_j^* - x_{\mathrm{lo},j}$. A zero gain row gives an infinite input bound. For every $0 < c \le c^*$ with $\Omega_c \subseteq \mathcal{N}$, the ellipsoid satisfies $\Omega_c \subseteq \mathcal{X}$, the control input $\boldsymbol{\pi}(\boldsymbol{x})$ lies within $\mathcal{U}$ on Ω_c, and Ω_c is forward invariant under the deployed loop; moreover, for $c < (s^*)^2$, every trajectory starting in Ω_c satisfies $\|\boldsymbol{x}_t - \boldsymbol{x}^*\|_{\mathbf{S}} \le \lambda(\sqrt{c})^t\|\boldsymbol{x}_0 - \boldsymbol{x}^*\|_{\mathbf{S}}$ with $\lambda(\sqrt{c}) = \hat{\vartheta}_{\mathbf{S}} + \eta_{\mathbf{S}} + \frac{1}{2}L_H^{\mathbf{S}}\sqrt{c} + \frac{1}{6}L_T^{\mathbf{S}}c < 1$.

Table 5: **Summary of benchmark dynamical systems.** n_x, n_u: state and control dimensions; Δt: discrete control period; $|u|_{\max}$: actuator bound per input channel; Transitions: total one-step transitions in the offline training dataset; Simulator: physics engine used for dynamics and ground-truth Jacobians.

System	n_x	n_u	Physical state coordinates	Δt (s)	$\lvert u\rvert_{\max}$	Transitions	Simulator
Pendulums							
Damped pendulum	2	1	$(\theta, \dot{\theta})$	0.02	8	2.1 M	exact (ODE)
Hanging double pendulum	4	2	$(\theta_{1,2}, \dot{\theta}_{1,2})$	0.01	6	2.1 M	exact (ODE)
Double pendulum (upright)	4	2	$(\theta_{1,2}, \dot{\theta}_{1,2})$	0.01	20	2.1 M	exact (ODE)
Cart-poles							
Cart-pole (ODE)	4	1	$(x, \theta, \dot{x}, \dot{\theta})$	0.02	10	2.1 M	exact (ODE)
Cart-pole (MuJoCo)	4	1	$(x, \theta, \dot{x}, \dot{\theta})$	0.02	10	2.1 M	exact (MuJoCo)
Cart-pole swing-up (MuJoCo)	5	1	$(x, \sin\theta, \cos\theta, \dot{x}, \dot{\theta})$	0.02	10	5.6 M	exact (MuJoCo)
Double cart-pole (MuJoCo)	6	1	$(x, \theta_{1,2}, \dot{x}, \dot{\theta}_{1,2})$	0.01	20	2.1 M	exact (MuJoCo)
Robot manipulators and legged systems							
Seven-joint arm (analytic)	14	7	joint positions, velocities	0.02	10	2.1 M	exact (ODE)
Kinova Gen3 arm (SPARK)	14	7	joint positions, velocities	0.02	10	50,000	exact (SPARK)
Unitree Go2 quadruped	35	12	joints, velocities, base pose	0.01	1.0	1.1 M	MuJoCo

Proof of Theorem 15 and Theorem 8. In Theorem 15, maximizing the linear forms $\hat{\mathbf{K}}_i \boldsymbol{e}$ and $\pm e_j$ over $\{\boldsymbol{e}^\mathsf{T}\mathbf{S}\boldsymbol{e} \le c\}$ gives $\sqrt{c\hat{\mathbf{K}}_i\mathbf{S}^{-1}\hat{\mathbf{K}}_i^\mathsf{T}} \le \bar{u}_i$ and $\sqrt{c(\mathbf{S}^{-1})_{jj}} \le b_j^{\pm}$ by the definition of c^* in (15), ensuring input and state constraint satisfaction on Ω_c. For any $\boldsymbol{x} \in \Omega_c \subseteq \mathcal{N}$ with $s = \|\boldsymbol{e}\|_\mathbf{S} \le \sqrt{c} \le s^*$, the triangle inequality and Lemma 14 yield

$$\|g_{\mathrm{loc}}(\boldsymbol{x}) - \boldsymbol{x}^*\|_\mathbf{S} \;\le\; \|\hat{\mathbf{F}}\|_\mathbf{S} s + \|\mathbf{F} - \hat{\mathbf{F}}\|_\mathbf{S} s + \|g_{\mathrm{loc}}(\boldsymbol{x}) - \boldsymbol{x}^* - \mathbf{F}\boldsymbol{e}\|_\mathbf{S} \;\le\; \lambda(s)\, s \;\le\; s,$$

proving forward invariance $g_{\mathrm{loc}}(\Omega_c) \subseteq \Omega_c$ and geometric decay when $c < (s^*)^2$. To obtain Theorem 8 (stated in Euclidean norms), note that for $\|\boldsymbol{e}\|_2 \le r$, $\|\boldsymbol{w}\|_\mathbf{S} \le \sqrt{\lambda_{\max}(\mathbf{S})}\|\boldsymbol{w}\|_2$ and $\|\boldsymbol{e}\|_2 \le \|\boldsymbol{e}\|_\mathbf{S}/\sqrt{\lambda_{\min}(\mathbf{S})}$ imply $\|g_{\mathrm{loc}}(\boldsymbol{x}) - \boldsymbol{x}^*\|_\mathbf{S} \le \big[\hat{\vartheta} + \sqrt{\kappa}(\eta + \frac{1}{2}L_H r + \frac{1}{6}L_T r^2)\big]\|\boldsymbol{e}\|_\mathbf{S}$. The bound in Theorem 8 makes this bracket at most 1 and ensures forward invariance on $\Omega_c \subseteq \{\|\boldsymbol{e}\|_2 \le r\}$. A strict bound gives exponential convergence. □

E Experimental Setup

A single export script generates the reported numerical results from evaluation logs. A separate pipeline reruns the experiments from trained checkpoints. Code, trained models, and simulated offline datasets will be released upon publication.

Benchmark systems and offline datasets. Table 5 summarizes the simulated dynamical systems studied in this paper. For the pendulums and cart-poles, the offline dataset consists of short trajectories starting from random states under random control inputs; for the torque-controlled seven-joint arm, a joint PD controller tracks random setpoints with exploration noise. On all benchmarks, the true system Jacobians $(\mathbf{A}_f, \mathbf{B}_f)$ for ground-truth evaluation are computed via central finite differences.

Network architectures and training regimes. Unless specified otherwise, the neural encoder Ψ_θ is a multilayer perceptron with two hidden layers of width 256 and ReLU activations. Standard Koopman and Koopman + invariance penalty each use 10 random seeds per system. The extended evaluation in Appendix F.1 includes additional regularization regimes with 5 seeds each: Gaussian state noise augmentation, a 1-Lipschitz encoder parameterized

with orthogonal layers (Wang and Manchester, 2023), and a penalty on the product of the latent feedback gain norm and encoder Lipschitz constant $\|\mathbf{K}_\Psi\|_2 L_\Psi$ (Abuduweili and Liu, 2024), evaluated with and without a variance floor to prevent feature collapse.

E.1 Training Objectives and Baseline Details

Unless stated otherwise, Koopman models are trained on $H = 15$-step windows from the offline trajectories. Each rollout starts from $\hat{\boldsymbol{z}}_0 = \Phi_\theta(\boldsymbol{x}_0)$ and advances according to $\hat{\boldsymbol{z}}_{k+1} = \mathbf{A}\hat{\boldsymbol{z}}_k + \mathbf{B}\boldsymbol{u}_k$. The discounted multi-step state prediction loss and the normalized latent linearity loss (Lusch et al., 2018) are

$$\mathcal{L}_{\text{pred}} = \frac{1}{\sum_k \gamma^k} \sum_{k=0}^{H-1} \gamma^k \left\| \mathbf{C}\hat{\boldsymbol{z}}_{k+1} - \boldsymbol{x}_{k+1} \right\|_2^2, \quad \mathcal{L}_{\text{lat}} = \frac{1}{\bar{\sigma}_\Psi^2 \sum_k \gamma^k} \sum_{k=0}^{H-1} \gamma^k \left\| [\hat{\boldsymbol{z}}_{k+1} - \Phi_\theta(\boldsymbol{x}_{k+1})]_\Psi \right\|_2^2, \tag{16}$$

where $[\cdot]_\Psi$ extracts the encoder coordinates (the last $N - n_x$ entries) of a latent vector and $\bar{\sigma}_\Psi^2$ is the batch variance of $\Psi_\theta(\boldsymbol{x})$, used to normalize the loss against feature scaling. The auxiliary inverse control loss (Abuduweili et al., 2026) is

$$\mathcal{L}_{\text{ctrl}} = \frac{1}{\sum_k \gamma^k} \sum_{k=0}^{H-1} \gamma^k \frac{1}{n_u} \left\| \mathbf{B}^+ \big(\Phi_\theta(\boldsymbol{x}_{k+1}) - \mathbf{A}\Phi_\theta(\boldsymbol{x}_k) \big) - \boldsymbol{u}_k \right\|_2^2. \tag{17}$$

Standard Koopman minimizes $\mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{ctrl}}$; Koopman + invariance penalty minimizes $\mathcal{L}_{\text{pred}} + \mathcal{L}_{\text{ctrl}} + w\,\mathcal{L}_{\text{lat}}$; and TIK fine-tuning in Appendix F.2 minimizes $\mathcal{L}_{\text{pred}}$ plus the complement-radius hinge penalty under (7). Standard training uses Adam with initial learning rate 10^{-3} decayed by 0.95 every 1000 steps, batch size 64, gradient norm clipping at 1.0, and 60,000 gradient steps with $\gamma = 0.8$. Swing-up models use $H = 40$ and $\gamma = 0.95$. TIK fine-tuning uses the schedule in Appendix F.2.

Neural dynamics baselines. The neural dynamics baseline models $\boldsymbol{x}_{t+1} = \boldsymbol{x}_t + h_\phi(\boldsymbol{x}_t, \boldsymbol{u}_t)$, where h_ϕ is an MLP with two hidden layers of width 256. It uses the same architecture, offline data, rollout horizon, and optimizer as the Koopman encoder. Neural model + NMPC linearizes $\boldsymbol{x} + h_\phi(\boldsymbol{x}, \boldsymbol{u})$ at the current state at every control step, centers the linearization at the target equilibrium, and solves the resulting convex QP with the same horizon, cost weights, and input bounds as Koopman MPC. Neural model + MPPI (Williams et al., 2017) uses the same stage and terminal cost, with sample count, exploration noise, and temperature tuned on the true linearization.

Dual-mode MPPI settings. In the swing-up study, each budget fixes the number of samples K and horizon H from measured per-step runtimes. For each budget, the terminal weight is selected on design starts disjoint from the evaluation starts. Only TIK's exploration scale $\sigma \in \{4, 6, 8, 12\}$ N and temperature $\lambda \in \{1, 3, 10, 30\}$ are tuned: they are selected on 64 design starts (24 hanging and 40 recovery starts across 4 models) by the highest success rate on the design starts. The baselines keep the default MPPI settings ($\sigma = 4$ N, $\lambda = 1$), so the comparison in Table 3 combines predictor quality with this tuning.

E.2 Local Linearization Estimator Validation

At any query point $(\boldsymbol{x}, \boldsymbol{u})$, we fit $\boldsymbol{x}_{i+1} \approx \boldsymbol{c} + \hat{\mathbf{A}}_f \boldsymbol{x}_i + \hat{\mathbf{B}}_f \boldsymbol{u}_i$ via distance-weighted least squares (6) on the $k = \max(200, 8(n_x + n_u + 1))$ nearest transitions in the offline dataset. Across the 129 trained checkpoints on the benchmark systems in Table 5, the estimated closed-loop spectral radius $\rho(\hat{\mathbf{F}})$ agrees with the finite-difference simulator reference $\rho(\mathbf{F})$ to within a median error of 0.0002 (95th percentile 0.0201, worst-case error at most 0.01).

E.3 Evaluation Protocol and Statistical Testing

For each controller, we evaluate $\rho(\mathbf{A}_{\mathrm{cl}})$, $\rho(\mathbf{F})$, and $\rho(\mathbf{\Pi A}_{\mathrm{cl}})$ at its closed-loop equilibrium. For uncentered controllers (Appendix F.1), equilibria of $g(\boldsymbol{x}) = f(\boldsymbol{x}, -\mathbf{K}\Phi_\theta(\boldsymbol{x}))$ are found via damped Newton iteration; an equilibrium is admissible if the control input is strictly unsaturated and the state lies within the 95th percentile of the training data distribution. Every admissible equilibrium lies at least 1.1×10^{-5} away from any ReLU activation boundary, ensuring that Φ_θ is continuously differentiable in a neighborhood of the equilibrium. Nonlinear verification rollouts perturb the equilibrium along the leading eigenvector of $\mathbf{F}$ and along four random directions and simulate the saturated closed loop for 3000 steps: a rollout is classified as stable if perturbations decay tenfold, unstable if they grow hundredfold, and undecided if $\rho(\mathbf{F})$ is too close to 1 to resolve within 3000 steps. For the recovery range in Figure 3, all pole angles of the initial state equal $\pm\theta$ and all other coordinates are zero. The saturated nonlinear loop runs for 1500 steps (2000 on the double cart-pole), and a run counts as recovered if the final state lies within 10^{-2} of the target. The tilts are 0.1 to 1.3 rad on the cart-poles, 0.01 to 0.25 rad on the double cart-pole, and 0.1 to 2.8 rad on the double pendulum. For each model, we average recovery over the two tilt signs. Each curve then averages these values across models. Where statistical significance is reported in the appendix, paired comparisons use the Wilcoxon signed-rank test for continuous metrics or the exact McNemar test for binary outcomes, with Holm step-down corrections across multiple comparisons.

F Supporting Experimental Results

F.1 Extended Analysis of the Nominal Certificate

Extended model classes. Table 6 extends Table 1 to a broader set of Koopman architectures pooled across the four open-loop unstable benchmark systems. Unlike in the main text, the controllers in Table 6, apart from the centered row, are evaluated without centering at the target equilibrium: each model is tested at the equilibrium $\boldsymbol{x}^*$ of its own deployed loop $g(\boldsymbol{x}^*) = \boldsymbol{x}^*$, provided an equilibrium exists within the actuator limits and training domain. A model is counted as *tested* if it has such an admissible equilibrium and its latent loop is nominally certified with $\rho(\mathbf{A}_{\mathrm{cl}}) < 1$.

In Table 6, 100% of tested Standard Koopman models and 100% of centered Standard Koopman models are unstable on the true system. Removing the auxiliary inverse control loss reduces the failure rate only slightly to 70%, with 28 of 40 models unstable and McNemar $p = 3.4 \times 10^{-3}$, showing that failures persist without the inverse control loss. We observe failures across all tested architectures that recover the physical state linearly from the latent state. These models have eigenvalues near the unit circle and LQR gains with norms up to $\|\mathbf{K}\|_2 = 6.3 \times 10^3$. Large gains can amplify errors in the input matching condition in (5). Autoencoders with a nonlinear MLP decoder have fewer local instabilities, but because the physical state $\boldsymbol{x} = g_\phi(\boldsymbol{z})$ is a nonlinear function of $\boldsymbol{z}$, state constraints and quadratic state costs become nonconvex in $\boldsymbol{z}$, so these models lose the convex QP structure of Koopman MPC. Table 7 breaks down the results by system across all pooled regularization regimes, including noise augmentation and encoder Lipschitz penalties: 73 of 115 trained controllers, or 66% of those admitting an unsaturated equilibrium, are unstable on the true system, whereas controllers on open-loop stable systems with $\rho(\mathbf{A}_f) < 1$ rarely fail.

Predicted versus deployed stability. Figure 2 shows the data behind the gain-scaling result of Section 6.1.

Table 6: **Nominal stability certificate failure rate across extended model classes** (four open-loop unstable systems pooled). Tested / trained: models that are nominally certified ($\rho(\mathbf{A}_{\text{cl}}) < 1$) and admit an unsaturated equilibrium within the training domain, out of all trained models. Unstable on true system (%): percentage of tested models whose deployed loop is physically unstable ($\rho(\mathbf{F}(\boldsymbol{x}^*)) > 1$).

Class	Model architecture	Tested / trained	Plant unstable (%) $\rho(\mathbf{F}(\boldsymbol{x}^*)) > 1$
Linear state recovery	Standard Koopman	39 / 40	100
	without inverse control loss	40 / 40	70
	centered at target equilibrium	40 / 40	100
	EDMD, quadratic polynomial dictionary	29 / 40	100
	EDMD, cubic polynomial dictionary	10 / 40	100
	EDMD, thin-plate radial basis functions	15 / 40	47
	EDMD, random Fourier features	0 / 40	no stabilizing LQR
	Autoencoder, linear state decoder	18 / 40	61
Published control methods	EDMD for Koopman MPC (Korda and Mezić, 2018)	23 / 40	61
	Deep Koopman with control (Shi and Meng, 2022)	40 / 40	28
	Stable operator (Mamakoukas et al., 2023)	40 / 40	100
	Stable embedding (Fan et al., 2026)	40 / 40	100
Nonlinear decoder	Autoencoder, nonlinear MLP decoder	37 / 40	5
	Autoencoder, latent linearity loss only	36 / 40	3
Ours	**TIK (ours)**	40 / 40	0

Table 7: **Breakdown of nominal certificate failures by benchmark system** (Standard Koopman and regularized training regimes pooled, excluding the invariance penalty; uncentered controllers). The first four systems are open-loop unstable ($\rho(\mathbf{A}_f) > 1$). Models with $\boldsymbol{x}^*$: models admitting an unsaturated equilibrium out of all trained models (all are nominally certified). Spectral radii and the physical linearization error ratio $\|\boldsymbol{\Delta}\|_2/d_1(\mathbf{F}_0)$ (Proposition 9) are medians across seeds.

	open-loop	models	certified	deployed, at $\boldsymbol{x}^*$		simulation	closure
System	$\rho(\mathbf{A}_f)$	with $\boldsymbol{x}^*$	$\rho(\mathbf{A}_{\text{cl}})$	$\rho(\mathbf{F}(\boldsymbol{x}^*))$	unstable	confirms	$\|\boldsymbol{\Delta}\|/d_1$
Double cart-pole (MuJoCo)	1.093	22/25	0.9935	1.0672	19 (76%)	21/21	4780
Cart-pole (ODE)	1.083	28/30	0.9863	1.0015	15 (50%)	27/27	18
Cart-pole (MuJoCo)	1.082	30/30	0.9840	0.9979	12 (40%)	30/30	27
Upright double pendulum	1.060	30/30	0.9769	1.0242	27 (90%)	30/30	19
Damped pendulum	0.990	20/20	0.9853	0.9889	5 (25%)	20/20	16
Seven-joint arm	0.980	18/18	0.8486	0.9930	3 (17%)	18/18	46
Hanging double pendulum	–	0/20		no admissible $\boldsymbol{x}^*$			

Insensitivity to LQR cost weights and centering. Sweeping the LQR control penalty $\mathbf{R}$ across three orders of magnitude over $[10^{-2}, 10^2]$ leaves 39 of 40 Standard Koopman controllers unstable at every weight. This occurs because the closed-loop error (5) is dominated by the open-loop state-transition error $\|\mathbf{J}_\theta \mathbf{A}_f - \mathbf{A}\mathbf{J}_\theta\|_2$ (median 182), which is independent of the feedback gain $\mathbf{K}$. By contrast, every TIK model remains stable across all tested LQR weights.

Soft invariance penalties. Adding the latent linearity loss $\mathcal{L}_{\text{lat}}$ with weight w in Table 8 eliminates the failure on the single cart-poles, where the learned encoder becomes nearly linear,

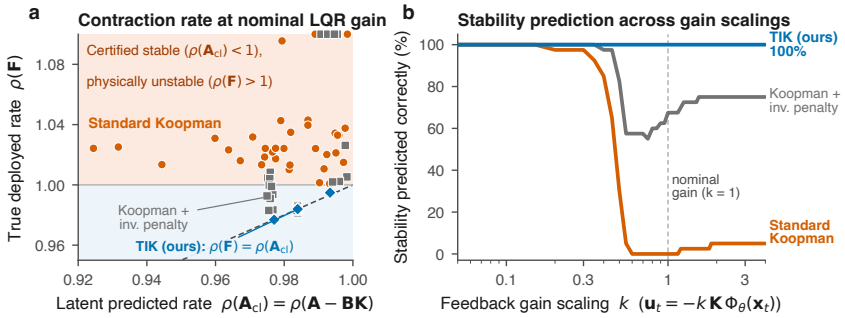


Figure 2: **Predicted and deployed spectral radii and stability as the feedback gain is scaled.** (a) Spectral radius predicted by the latent model, $\rho(\mathbf{A}_{\mathrm{cl}}) = \rho(\mathbf{A} - \mathbf{BK})$, plotted against the deployed spectral radius $\rho(\mathbf{F})$ at the nominal LQR gain for Standard Koopman (orange circles), Koopman + invariance penalty (gray squares), and TIK (blue diamonds). (b) Percentage of models (40 seeds per method) whose latent stability prediction matches the true deployed stability as the feedback gain is scaled by $k \in [0.05, 4]$.

Table 8: **Effect of the soft invariance penalty weight w.** Values are medians over n seeds. Setting $w = 0$ corresponds to Standard Koopman. Latent and state errors are one-step prediction errors on held-out validation data; nonlin. measures the relative variation of the encoder Jacobian across the dataset (0 for an affine encoder). The final column is the number unstable among models with an admissible equilibrium. The sweep includes runs evaluated at their own equilibria, so these counts differ from the centered comparison in Table 1.

System	w	n	latent err.	state err.	14-step err.	nonlin.	$\rho(\mathbf{A}_{\mathrm{cl}})$	$\rho(\mathbf{F}(\boldsymbol{x}^*))$	$\|\mathcal{E}\|_2$	unstable
Cart-pole (ODE)	0	10	0.526	0.203	0.013	0.006	0.9873	1.0199	134	10/10
Cart-pole (ODE)	0.1	10	0.002	0.061	0.013	0.006	0.9839	0.9842	0.34	0/10
Cart-pole (ODE)	1	10	0.001	0.035	0.012	0.003	0.9839	0.9846	0.21	0/10
Cart-pole (ODE)	10	10	0.000	0.027	0.012	0.001	0.9839	0.9843	0.09	0/10
Cart-pole (MuJoCo)	0	10	0.528	0.217	0.012	0.006	0.9886	1.0210	131	10/10
Cart-pole (MuJoCo)	0.1	10	0.002	0.057	0.012	0.006	0.9840	0.9850	0.33	0/10
Cart-pole (MuJoCo)	1	10	0.001	0.033	0.011	0.003	0.9839	0.9833	0.18	0/10
Cart-pole (MuJoCo)	10	10	0.000	0.028	0.011	0.001	0.9839	0.9838	0.076	0/10
Double cart-pole	0	10	0.525	0.457	0.391	0.091	0.9918	1.1480	265	9/9
Double cart-pole	1	10	0.007	0.285	0.377	0.069	0.9944	1.1311	34	5/6
Double cart-pole	10	5	0.004	0.261	0.379	0.043	0.9929	1.0059	11	5/5
Seven-joint arm	0	5	0.620	0.050	0.062	0.079	0.9969	0.9992	14	1/5
Seven-joint arm	1	5	0.003	0.047	0.064	0.116	0.9816	0.9777	0.056	0/5
Seven-joint arm	10	5	0.002	0.056	0.068	0.079	0.9825	0.9779	0.082	0/5

but leaves 5 of 6 double cart-pole models unstable at $w = 1$. Minimizing the average latent prediction error on the training data does not enforce the derivative matching conditions sufficient for Theorem 2. Directly penalizing the error Jacobian $\|\mathcal{E}\|_2$ at the target during training still leaves 45 of 46 runs unstable, and post-hoc anchored invariance (12) converges on 0 of 9 uncentered double cart-pole models.

Predictive safety filters and unknown equilibria. Wrapping a Standard Koopman controller in a predictive safety filter (Wabersich and Zeilinger, 2018) that predicts trajectories using the learned latent model still violates state constraints in 41% of MuJoCo cart-pole episodes, compared to 0% when the safety filter predicts using the local linear fit (6), with Holm $p = 0.020$. When the equilibrium of an uncentered controller is unknown, estimating

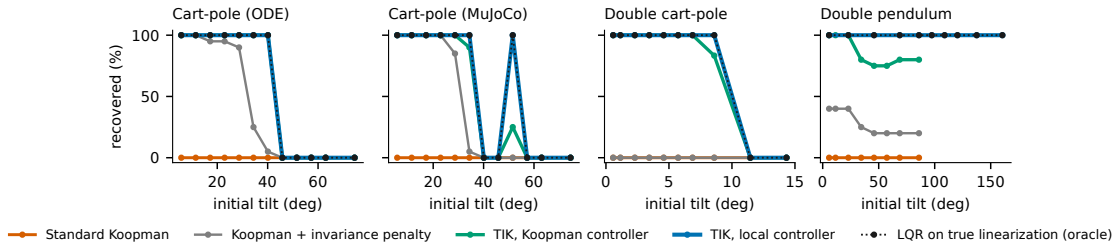


Figure 3: **Empirical recovery from initial tilts.** Fraction of runs returning to the upright with all poles tilted at rest (two tilt signs, up to 10 seeds per learned curve, one oracle). TIK local feedback uses the fitted physical LQR gain $\hat{\mathbf{K}}$. The lifted TIK controller uses the checkpoint after additional closed-loop training (Appendix F.2).

it from 400 offline transitions and applying the offline stability test yields correct verdicts for 43 of 105 Standard Koopman models with zero false conclusions across safety factors in $[1, 3]$ and neighbor counts k.

F.2 Training and Certifying TIK

Algorithm 1 summarizes the complete TIK procedure. Each TIK model is fine-tuned from the trained Standard Koopman checkpoint of the same seed using Adam with learning rate 10^{-4}, initializing $\mathbf{W}$ to the trained $\mathbf{A}$ and $\mathbf{N}_j = \mathbf{0}$, and using a soft hinge penalty $100 \max(0, \rho_c - \bar{\rho})^2$ on the spectral radius ρ_c of the orthogonal complement block $(\mathbf{I} - \mathbf{J}_\theta \mathbf{J}_\theta^+)\mathbf{W}(\mathbf{I} - \mathbf{J}_\theta \mathbf{J}_\theta^+)$ to encourage the dynamics on the complement to be Schur. Proposition 4(b) requires a separate Schur check. Regulation models are fine-tuned for 6000 steps ($\bar{\rho} = 0.98$), and bilinear swing-up models for 20,000 steps ($\bar{\rho} = 0.999$). On the double cart-pole, where the deployed spectral radius of Standard Koopman ranges from 1.002 to 113, fine-tuning with TIK brings every model below $\rho(\mathbf{F}) \leq 0.9949 < 1$. Independent replications using freshly drawn offline datasets yield 10/10 stable models on the double cart-pole, 6/6 at an off-origin equilibrium, and 12/12 on the upright double pendulum.

Closed-loop training extends the recovery range. Because the offline training trajectories do not visit all large-tilt closed-loop states, we fine-tune TIK in closed-loop rounds (4 rounds with 6,000 updates on the cart-poles, and 20 rounds with 40,000 updates and 70% replay-pool sampling on the double cart-pole and double pendulum). In each round, we simulate the current Koopman controller from 300 initial states (all pole angles tilted by up to $\theta_{\max}$ with 10% actuator exploration noise) for 250 steps, and add 15-step windows with pole angles below $3\theta_{\max}$ to a replay pool ($\theta_{\max} = 0.7$ rad on the cart-poles, 0.15 rad on the double cart-pole, and 1.5 rad on the double pendulum). Table 9 reports the result: tangent invariance is preserved by construction, and all 39 models that complete training are stable on the true system (one double cart-pole seed diverged to NaN and is excluded). With the same closed-loop data but no tangent matching, only 2 of 30 controllers are stable on the plant.

Checking the offline error bounds. In Table 2, the two-scale protocol fits a local cubic polynomial (Corollary 12) on 200000 transitions within radius 5×10^{-2} of the target and fits the linear model (6) on 50000 transitions within radius 5×10^{-3}. Table 10 compares the estimated bounds with simulator derivatives across the fits in Table 2 and additional noise regimes. The error bound (8) covers the simulator Jacobian error in every listed fit. However,

Table 9: **Recovery range of the TIK controller before and after training on closed-loop rollouts.** Largest tilt (rad) with at least half of the runs recovered, median over models; "not worse": models whose range did not shrink; "control": (model, tilt) cells recovered by the model trained with the same data but no invariance constraint.

System	Oracle	Before ↑	After ↑	Not worse ↑	Control cells ↓
Cart-pole (ODE)	0.70	0.60	0.70	10/10	0
Cart-pole (MuJoCo)	0.90	0.75	0.75	6/10	6
Double cart-pole	0.15	0.06	0.15	9/9	0
Double pendulum	>1.5	0.40	1.50	10/10	–

Table 10: **Offline error bounds checked against simulator derivatives.** The study includes the regimes in Table 2 and additional noise levels (100 fits per cart-pole row and 40 per double cart-pole row).

System	Offline data regime	Noise ϵ	Bound η contains true error ↑	Curvature bound holds ↑
Cart-pole swing-up	Near the target	0	100/100	80/100
	Two scales	0	100/100	100/100
	Two scales	10^{-5}	100/100	100/100
	Two scales, noisy states	10^{-7}	100/100	100/100
	Closed-loop trajectories	3×10^{-8}	100/100	100/100
Double cart-pole	Two scales	0	40/40	40/40
	Two scales	10^{-6}	40/40	40/40
	Two scales, noisy states	10^{-9}	40/40	40/40
	Closed-loop trajectories	10^{-10}	40/40	40/40

the curvature estimate underestimates simulator curvature in some fits, so the checks do not establish all assumptions of Theorem 6.

Local MPC implementation. In condensed form, the Koopman MPC quadratic program (2) optimizes over the $N_h = 20$ control inputs and is solved via DAQP in a median of $7.3\,\mu$s; near the target equilibrium, its effective feedback gain agrees with the infinite-horizon Koopman LQR gain $\mathbf{K}$ to within 10^{-12} on every model (Proposition 7).

F.3 Swing-Up Ablation, Safe-Control-Gym, and Manipulator Details

Swing-up ablation and equal-effort comparison. On the MuJoCo cart-pole swing-up task with continuous angle embedding $(x, \sin\theta, \cos\theta, \dot{x}, \dot{\theta})$, an episode succeeds if the pole is held upright ($\cos\theta > 0.98$) over the final 2 s of a 10 s rollout without touching the rail stops. A single MPPI planning step with 512 trajectory samples requires a median of 10.3 ms on the bilinear Koopman model versus 17.4 ms on the neural network model; under tighter time budgets, each planner selects the sample count and horizon that fits the budget and maximizes validation success. Table 11 presents the component ablation discussed in Section 6.3, illustrating the separate effects of tangent matching and bilinear actuation. When both TIK and Neural model + MPPI are trained from scratch for 3× as many gradient steps (180,000 steps), the neural network recovers from aggressive kicks slightly more often when compute is unconstrained (100% vs. 92%), whereas TIK matches its 100% swing-up rate, achieves 99% combined success at 1 ms (vs. 58% for Neural MPPI), and remains the only controller capable of executing within the 0.25 ms budget (80% success).

Table 11: **Component ablation on cart-pole swing-up** (20 models per row). Entry is the fraction of MPPI-only runs reaching the upright neighborhood. Success is the fraction completing the task with MPPI alone. The two MPPI + LQR columns report overall success after switching to lifted LQR or LQR on the fitted physical linearization. The final column counts models with an unstable deployed lifted LQR loop among those with an admissible upright fixed point.

			MPPI only (%)		MPPI + LQR (%)		
Model variant	Matched	Bilinear	Entry	Success	Lifted	Fitted	Unstable
Standard Koopman	–	–	95	0	0	91	20/20
Fine-tuned without TIK	–	–	74	0	0	76	20/20
Fine-tuned without TIK	–	✓	100	57	0	99	20/20
TIK, linear input	✓	–	41	6	11	16	0/20
TIK (ours)	✓	✓	100	95	88	99	0/20

Table 12: **Real-time MPC on simulated seven-joint manipulators under per-step computation budgets.** Success rate (%) on Kinova Gen3 reaching around an obstacle ($10\times$ 30 episodes) and on a torque-controlled arm reaching six goals (5×36 episodes). The linear MPC reference uses a fit from logged transitions on Kinova and a single workspace-center linearization on the six-goal arm. The archived protocol gives the reference configuration for this latter baseline and the Go2 baseline, but leaves their Jacobian source unspecified.

		Kinova reach (%) ↑		Six-goal arm (%) ↑
Class	Method	5 ms	20 ms	1 ms
Linear	Linear MPC	100	100	100
Neural	Neural model + MPPI	1	91	0
	Neural model + NMPC	–	–	92
Koopman	Standard Koopman	100	100	0
Ours	**TIK (ours)**	100	100	100

Why the two LQR controllers can differ. In centered coordinates, their unsaturated actions satisfy

$$\boldsymbol{u}_{\text{lift}} - \boldsymbol{u}_{\text{fit}} = -\mathbf{K}\big(\Phi_\theta(\boldsymbol{x}) - \mathbf{J}_\theta \boldsymbol{x}\big) = o(\|\boldsymbol{x}\|_2).$$

Thus Proposition 4 gives equal feedback derivatives at the target, not equal actions throughout the upright neighborhood. The 88% versus 99% success rates in Table 11 are compatible with this local result. They do not isolate the cause of the gap: that would require matched switching states, thresholds, and saturation rules. The budget sweep in Table 3 is a separate protocol, even where a reported percentage is the same.

Safe-control-gym and Kinova Gen3 details. Table 13 reports the complete results on the safe-control-gym cart-pole stabilization benchmark (Yuan et al., 2022). TIK reduces state constraint violations from 96% for Standard Koopman to 4%, matching LQR on the oracle linearization (2%). On the Kinova Gen3 arm (50 Hz joint-acceleration control reaching 10 goal poses around a spherical obstacle, solved via OSQP), TIK solves each QP step in a median of 2.1 ms, reaches the goal in 100% of episodes at 5 ms (Holm $p = 3.1 \times 10^{-5}$ vs. Neural model + MPPI), and achieves lower closed-loop trajectory cost at 20 ms (543 vs. 566 for Standard Koopman, $p = 1.6 \times 10^{-4}$).

Table 13: **Safe-control-gym cart-pole stabilization benchmark.** Percentage of episodes that violate state constraints (↓) and that successfully stabilize to the target (↑) across 10 models × 25 initial states.

Class	Controller	Constraint violations (%) ↓	Stabilized (%) ↑
Linear	Oracle LQR (simulator Jacobians)	2	100
	Fitted linear LQR (offline logs)	2	100
	Oracle linear MPC (benchmark)	32	100
Neural	Neural model + NMPC	28	91
	Neural model + MPPI	7	0
Koopman	Standard Koopman	96	0
	+ safety filter on latent model	86	0
	+ safety filter on linear fit	72	0
Ours	**TIK (ours)**	4	100

F.4 Multiple Equilibria and Higher-Dimensional Robotic Systems

Proposition 16 (Obstruction to multi-equilibrium invariance with a linear input)**.** *Consider a linear latent model* $\hat{\boldsymbol{z}}_{t+1} = \mathbf{A}\boldsymbol{z}_t + \mathbf{B}\boldsymbol{u}_t$ *with state recovery* $\mathbf{C}\mathbf{J}_i = \mathbf{I}$ *at equilibria* $i = 1, \dots, k$*. If* $\mathbf{A}\mathbf{J}_i = \mathbf{J}_i\hat{\mathbf{A}}_i$ *and* $\mathbf{B} = \mathbf{J}_i\hat{\mathbf{B}}_i$ *for all* $i = 1, \dots, k$*, then: (a) the Markov parameters satisfy* $\mathbf{C}\mathbf{A}^m\mathbf{B} = \hat{\mathbf{A}}_i^m\hat{\mathbf{B}}_i$ *for every* $m \geq 0$ *and every* i*; and (b) if each pair* $(\hat{\mathbf{A}}_i, \hat{\mathbf{B}}_i)$ *is controllable, then* $\hat{\mathbf{A}}_i = \hat{\mathbf{A}}_1$ *and* $\hat{\mathbf{B}}_i = \hat{\mathbf{B}}_1$ *for all* i*.*

Proof. (a) By induction on m, $\mathbf{A}^m\mathbf{J}_i = \mathbf{J}_i\hat{\mathbf{A}}_i^m$, so $\mathbf{C}\mathbf{A}^m\mathbf{B} = \mathbf{C}\mathbf{A}^m\mathbf{J}_i\hat{\mathbf{B}}_i = \mathbf{C}\mathbf{J}_i\hat{\mathbf{A}}_i^m\hat{\mathbf{B}}_i = \hat{\mathbf{A}}_i^m\hat{\mathbf{B}}_i$. (b) By part (a), the matrix $\mathcal{C} = \mathbf{C}[\mathbf{B}, \mathbf{A}\mathbf{B}, \dots, \mathbf{A}^{n_x-1}\mathbf{B}]$ equals the controllability matrix of every pair $(\hat{\mathbf{A}}_i, \hat{\mathbf{B}}_i)$ and therefore has full row rank n_x. Since $\hat{\mathbf{A}}_i\mathcal{C} = \mathbf{C}[\mathbf{A}\mathbf{B}, \dots, \mathbf{A}^{n_x}\mathbf{B}]$ is independent of i, right-multiplying by $\mathcal{C}^+$ gives $\hat{\mathbf{A}}_i = \hat{\mathbf{A}}_1$, and $\hat{\mathbf{B}}_i = \mathbf{C}\mathbf{B} = \hat{\mathbf{B}}_1$. □

Corollary 17 (Non-existence of exact linear-input liftings across distinct equilibria)**.** *Let* Φ *be a finite-dimensional differentiable lifting with linear state recovery* $\mathbf{C}\Phi(\boldsymbol{x}) = \boldsymbol{x}$*, and let* $\mathbf{A}, \mathbf{B}, \mathbf{C}$ *be constant matrices. If* $\Phi(f(\boldsymbol{x}, \boldsymbol{u})) = \mathbf{A}\Phi(\boldsymbol{x}) + \mathbf{B}\boldsymbol{u}$ *holds for all* $(\boldsymbol{x}, \boldsymbol{u})$*, then all equilibria of* f *with controllable linearizations share the same linearization* $(\mathbf{A}_f, \mathbf{B}_f)$*. In particular, no such exact lifting exists for a pendulum, whose downward and upright equilibria have distinct controllable linearizations.*

Proof. Differentiating $\Phi(f(\boldsymbol{x}, \boldsymbol{u})) = \mathbf{A}\Phi(\boldsymbol{x}) + \mathbf{B}\boldsymbol{u}$ at any equilibrium $(\boldsymbol{x}_e, \boldsymbol{u}_e)$ yields the relations $\mathbf{J}_\Phi(\boldsymbol{x}_e)\mathbf{A}_f(\boldsymbol{x}_e) = \mathbf{A}\mathbf{J}_\Phi(\boldsymbol{x}_e)$ and $\mathbf{J}_\Phi(\boldsymbol{x}_e)\mathbf{B}_f(\boldsymbol{x}_e) = \mathbf{B}$ with $\mathbf{C}\mathbf{J}_\Phi(\boldsymbol{x}_e) = \mathbf{I}$, so Proposition 16(b) applies directly. □

Bilinear input allows the effective matrices to vary across goals. The restriction in Proposition 16 therefore no longer follows. Joint feasibility still depends on the goal embeddings and fitted Jacobians. See also Ge et al. (2026), Liu et al. (2025), Shang et al. (2026b) for related restrictions across operating points. At goal i, write $\boldsymbol{z}_i = \Phi_\theta(\boldsymbol{x}_i)$, $\mathbf{A}_i^z = \mathbf{A} + \sum_j u_{ij}^*\mathbf{N}_j$, and $\mathbf{B}_i^z = \mathbf{B} + [\mathbf{N}_1\boldsymbol{z}_i \ \cdots \ \mathbf{N}_{n_u}\boldsymbol{z}_i]$. Training projects onto the local matching equations $\mathbf{A}_i^z\mathbf{J}_i = \mathbf{J}_i\hat{\mathbf{A}}_i$ and $\mathbf{B}_i^z = \mathbf{J}_i\hat{\mathbf{B}}_i$ and constrains each $\boldsymbol{z}_i$ to be fixed under $\boldsymbol{u}_i^*$. At fixed encoder weights these constraints are linear in $\mathbf{A}, \mathbf{B}, \{\mathbf{N}_j\}$, and we check feasibility numerically. For a local model in goal deviations, exact matching and the Schur, stabilizability, and detectability conditions of Proposition 4(b) give the fitted physical LQR gain on the tangent space.

Seven-joint arm across compute budgets. On the torque-controlled seven-joint manipulator with a 14-dimensional state (Table 5), the 6 goal poses have true linearizations differing by up to 0.085. With the bilinear TIK model, multi-point tangent invariance holds

Table 14: **Torque-controlled seven-joint manipulator reaching six goals across computation budgets.** Success rate (%) over 5 models and 36 episodes per model per budget. TIK's 10 ms run had deadline misses under machine load. A repeat at lower load reached 100%. The table reports the original budget sweep.

		Success rate (%) at per-step compute budget ↑				
Class	Controller	1 ms	2 ms	5 ms	10 ms	20 ms
Linear	Linear MPC (workspace center)	100	100	100	100	100
Neural	Neural model + MPPI	0	1	98	98	98
	Neural model + NMPC (relinearized)	92	100	100	97	100
Koopman	Standard Koopman MPC	0	0	0	0	0
Ours	**TIK (ours, multi-point QP)**	100	100	100	89	97

Table 15: **Unitree Go2 quadruped (35D state, 12D input) posture recovery under pushes and domain shifts** (150 episodes per entry).

		Success rate (%) ↑						
Class	Method	Nominal	Noise	Friction	0.8× mass	1.2× mass	Falls (%) ↓	ms/step ↓
Linear	Linear MPC (single standing pose)	3	–	–	–	–	50	0.68
Neural	Neural model + MPPI + filter	47	–	–	–	–	15	6.74
	Neural model + NMPC	51	–	–	–	–	4	5.35
Koopman	Standard Koopman QP	26	–	–	–	–	10	0.01
Ours	**TIK (ours, multi-point QP)**	67	67	73	57	74	21	0.83

to within 10^{-12} across all 5 seeds, stabilizing all 30/30 goal–model pairs, compared with 2/30 for Standard Koopman. The offline test returned stable verdicts for 29/30 TIK pairs and 0/30 Standard Koopman pairs, with no false stable verdicts in simulator checks. In Table 14, Standard Koopman fails at all budgets (0%) and Neural MPPI fails at 1–2 ms (0%–1%), whereas multi-point TIK achieves 100% at 1–5 ms and outperforms Neural NMPC at 1 ms (92%, Holm $p = 0.029$). Linear MPC at the workspace center also reaches 100% at 1 ms, so this task shows no success-rate advantage for TIK over that baseline.

Unitree Go2 quadruped posture recovery and contact boundaries. We evaluate a Unitree Go2 quadruped with a 35-dimensional state in MuJoCo Menagerie under 100 Hz joint-target control over a 500 Hz PD loop, holding 7 standing postures and recovering from impulse pushes without foot slip in Table 15. Across the 7 postures, the true linearized state matrices differ by 5–21%. Linear MPC designed at a single pose succeeds in 3% of episodes and the robot falls in 50%. Standard Koopman succeeds in 26%. Multi-point TIK with a smooth encoder, goal offsets consistent with the model, and tightened constraints reaches 67% under nominal conditions and 67%–74% under sensor noise, friction changes, and +20% payload mass, outperforming Neural MPPI at 47% ($p = 2.3 \times 10^{-3}$) and Neural NMPC at 51% ($p = 9.5 \times 10^{-3}$) while running 6.4–8.1× faster at 0.83 ms per step. A Koopman model without tangent matching, using the same smooth encoder and constraint tightening, reaches 77%. The offline test is inconclusive in all 150 cases, with $r_{\mathbb{R}}(\hat{\mathbf{F}}) \leq 4.4 \times 10^{-4} < \eta$. The deployed spectral radius is near one (median 1.0000), and nonsmooth contact transitions can violate the local differentiability assumptions.

F.5 Dual-Mode QP-MPC on the Double Cart-Pole

We evaluate the dual-mode QP-MPC of Section 5.3 on the double cart-pole from tilted starts: $|\theta_1| \in [18°, 28°]$, θ_2 of opposite sign, and cart offsets up to 0.3 m (20 starts per model). The

local controller is $\boldsymbol{u} = -\hat{\mathbf{K}}\boldsymbol{x}$. With only the original offline data, TIK has one-step error 0.08 on local-controller trajectories and 4% MPC success. We then fine-tune for 80,000 updates, drawing 60% of each batch from 15-step windows under the local controller from tilted starts with input noise on disjoint seeds, and adding near-target windows. We save checkpoints every 5,000 updates and select by MPC success on disjoint validation starts.

The selected models have one-step error 0.01 and succeed in 83% of rollouts across 20 models (83% and 82% in the two regimes; up to 95% for one seed). Of all rollouts, 84% enter Ω_c; every model succeeds at least once. Fitted linear MPC with the same terminal set reaches 90%, and fitted LQR applied from the start reaches 90%. Standard Koopman dual-mode MPC has no successful rollout (10/10 models fail). Training on the same closed-loop data without tangent matching gives 16% success (8/10 models never succeed).