Journal Title
XX(X):1–24
©The Author(s) 0000
Reprints and permission:
sagepub.co.uk/journalsPermissions.nav
DOI: 10.1177/ToBeAssigned
www.sagepub.com/
SAGE

SACK: Safe Active Continual Koopman Learning for Uncertain Systems with Contractive Guarantees

Chandan Kumar Sah*, Rajpal Singh* and Jishnu Keshavan

Abstract

Koopman operator theory provides a powerful framework for representing nonlinear dynamics through a linear operator acting on lifted observables, enabling the application of linear control techniques to nonlinear systems. However, Koopman models are typically learned offline and often degrade under model uncertainty and distributional shifts between training and deployment. Existing adaptive Koopman methods either incur significant computational overhead or lack formal safety guarantees, limiting their applicability to real-time, safety-critical robotic systems. We propose SACK, a unified framework for safe continual Koopman learning that enables efficient online adaptation of learned models during task execution. A nominal Koopman model is first learned offline and subsequently refined through a contractive adaptation law with theoretical convergence guarantees. To improve data efficiency, the framework integrates an active learning strategy that acquires informative data while accomplishing task objectives. We derive explicit bounds on model uncertainty and incorporate them within a robust Model Predictive Control framework to establish formal safety guarantees. To reduce conservatism, we further introduce a conformal prediction-based tightening mechanism that calibrates safety margins online from observed residuals. Extensive simulations and hardware experiments across diverse robotic platforms demonstrate superior performance compared with state-of-the-art methods in safety-critical environments.

Keywords

Koopman Operator, Online Adaptation, Nonlinear Systems, Model Predictive Control

1 Introduction

Control of nonlinear dynamical systems remains a central challenge in engineering, arising across robotic manipulation, aerial vehicles, legged locomotion, and a wide range of physical and biological systems (Dixon et al. 2003; Rasmussen et al. 2011; Wang et al. 2008). Although model-based control techniques for nonlinear systems are well developed, they are often system-specific and computationally demanding, particularly when optimality and hard constraints must be considered (Strogatz 2018). In contrast, linear control offers scalable and computationally efficient tools with well-established theoretical guarantees (Ogata et al. 2010). This disparity has motivated the development of representations that retain the expressive capability required for nonlinear systems while enabling the use of linear analysis and control techniques.

Koopman operator theory provides a principled framework for obtaining such representations by describing nonlinear dynamics through a linear operator acting on nonlinear observable functions (Koopman 1931; Korda and Mezić 2018a). Since an exact Koopman representation is generally infinite dimensional, practical implementations rely on finite-dimensional approximations learned from data using methods such as Dynamic Mode Decomposition, Extended Dynamic Mode Decomposition, and neural-network-based lifting functions (Champion et al. 2019; Li et al. 2017; Lusch et al. 2018). These models have been successfully employed for the modeling and control of robotic manipulators (Sah et al. 2024; Singh et al. 2025b), soft robots (Bruder et al. 2020, 2025), aerial systems (Zinage and Bakolas 2021; Folkestad et al. 2022; Manaa et al. 2024), and legged robots (Krolicki et al. 2022; Yang and Bhounsule 2025). More recently, Koopman representations have also been incorporated into safety-critical control architectures, where their lifted linear structure enables tractable prediction and safety verification (Folkestad et al. 2020).

Despite these advances, Koopman models remain fundamentally data-driven and therefore inherit the limitations of the data on which they are trained. Changes in payload, actuator characteristics, environmental disturbances, contact interactions, or other operating conditions can cause the deployed dynamics to depart substantially from those represented during training, rendering an offline-learned model inaccurate. Online adaptation (Singh et al. 2025a; Uchida and Duraisamy 2024) offers a natural remedy, but its effectiveness is fundamentally limited by the data available during operation. In particular, tracking controllers tend to generate smooth and correlated trajectories that provide poor excitation for identifying changing dynamics. Active learning and information-seeking control address this

The authors are with the Department of Mechanical Engineering, Indian Institute of Science, Bangalore, Karnataka 560012, India
*These authors have contributed equally to the present work.

Corresponding author:
Rajpal Singh, Department of Mechanical Engineering, Indian Institute of Science, Bangalore, Karnataka 560012, India.
Email: rajpalsingh@iisc.ac.in

limitation by deliberately shaping the system trajectory to acquire informative measurements (Abraham and Murphey 2019; Soloperto et al. 2020; Baltussen et al. 2025). However, such excitation creates a competing risk. The actions required for learning may degrade task performance or violate state and input constraints. More importantly, when safety constraints are enforced using an inaccurate or continually adapting model, constraint satisfaction can no longer be reliably guaranteed unless the associated model uncertainty is explicitly accounted for. Thus, the controller must not only generate informative trajectories, but also ensure that exploration and model updates do not compromise constraint satisfaction during adaptation. This tension is most acute immediately after a distributional shift, when additional information is most valuable but the model used to generate safe exploratory actions is least reliable. The central challenge is therefore how a robot can effectively adapt its learned dynamics under distributional shift without sacrificing safety during closed-loop operation.

1.1 Contributions

To address this problem, we propose *Safe Active Continual Koopman Learning (SACK)*, a framework that couples continual Koopman adaptation with active data acquisition and safety-aware model predictive control. A nominal Koopman representation is first learned offline from data collected under nominal operating conditions. During deployment, the lifting map is retained while the Koopman operator matrices are updated recursively using closed-loop measurements. In contrast to iterative neural-network retraining (Singh et al. 2025a; Uchida and Duraisamy 2024), the proposed update admits a closed-form implementation and yields parameter-error dynamics whose contraction is governed explicitly by the windowed regressor Gramian. This relationship provides an analytical connection between the informativeness of the collected data and the rate at which the learned model adapts.

To exploit this connection, we incorporate a D-optimal information criterion within a receding-horizon controller that actively shapes future state-input trajectories to improve the conditioning of the regressor Gramian while simultaneously accomplishing the control task. Since such information-seeking behavior is performed using an uncertain and continually changing model, we additionally propagate the adaptation error into tightened control-barrier-function constraints within the MPC. The resulting analytical bounds provide deterministic safety margins, while a conformal residual-calibration mechanism is introduced to reduce the conservatism associated with worst-case uncertainty bounds. Consequently, SACK provides a closed loop framework that uses the information content of closed-loop data to govern model adaptation, while explicitly accounting for the resulting model uncertainty in the safety constraints.

The main contributions of this work are:

1. **Contractive continual Koopman adaptation.** We derive a closed-form recursive update law for the Koopman operator and establish exponential contraction of the parameter estimation error under persistent excitation. For time-varying dynamics, we derive an explicit ultimate bound and characterize the effects of the adaptation window and exponential forgetting on the achievable estimation accuracy.
2. **Active data acquisition coupled to adaptation.** We formulate a D-optimal active-learning objective based on the predicted regressor Gramian and integrate it within receding-horizon control. We further derive a lower bound on the minimum eigenvalue of the predicted Gramian, establishing an explicit connection between the information-seeking objective and the excitation required for contraction of the adaptive model.
3. **Safety-aware continual learning.** We derive model-error bounds from the adaptation dynamics and incorporate them within tightened CBF constraints in the MPC. To reduce the conservatism of deterministic worst-case bounds, we further introduce conformal residual calibration for constructing data-driven safety margins with probabilistic coverage guarantees.
4. **Validation across diverse robotic systems.** We evaluate SACK on a 3R manipulator and planar quadrotors under large parameter variations and external disturbances, a 7-DoF Franka Research 3 manipulator in high-fidelity simulation, and a TurtleBot3 mobile robot in hardware with deliberately induced model mismatch. Detailed comparisons with adaptive, active-learning, and non-adaptive alternatives demonstrate the efficacy of the proposed framework under distributional shift and safety constraints.

The remainder of the paper is organized as follows. Section 2 discusses related work on adaptive Koopman learning, active model learning, and safe learning under model uncertainty. Section 3 presents the adaptive Koopman formulation and its convergence analysis. Section 4 develops the active-learning MPC formulation and excitation guarantees. Sections 5 and 6 present the analytical and data-driven safety formulations, respectively. Section 7 presents the simulation and experimental results, followed by conclusions in Section 8.

2 Related Work

The problem considered in this work lies at the intersection of online Koopman adaptation, active model learning, and safety-aware control under model uncertainty. We review these directions with emphasis on methods that refine learned dynamics during deployment.

2.1 Adaptive and Continual Koopman Learning

Most Koopman-based control methods rely on models learned offline under nominal operating conditions (Sah et al. 2024; Singh et al. 2025b; Folkestad et al. 2022; Bruder et al. 2025; Yang and Bhounsule 2025). Although such models have demonstrated strong performance, their accuracy can deteriorate when the deployed dynamics depart from the training distribution. These challenges motivate the development of adaptive Koopman frameworks that refine the learned model online during deployment.

As such several recent works have been proposed to address online Koopman adaptation. The study in Singh

et al. (2025a) augments an offline model with an auxiliary network that learns the residual mismatch in real time. Uchida and Duraisamy (2024) extend this approach to jointly optimize the lifting map and operator matrices online using soft target-network stabilization. The study in Li et al. (2024) incrementally expands the lifted-state dimension to provably reduce approximation error for high-dimensional legged robots, while a Bayesian meta-learning prior over Koopman operators is adopted in Selim et al. for closed-form, optimization-free adaptation. Each, however, carries a critical limitation. The neural updates of Singh et al. (2025a); Uchida and Duraisamy (2024) require iterative gradient steps at every control step, forcing a fundamental trade-off: the number of gradient steps must be small enough to meet latency requirements, yet sufficient to reduce model error meaningfully. Under rapid distributional shift, precisely the very conditions that motivate adaptation, this trade-off typically favors speed over accuracy. The incremental scheme of Li et al. (2024) sidesteps convergence by growing the model, but scales the lifted-state MPC, adding to real-time computational burden. In Selim et al., closed-form updates are obtained only when the deployment distribution lies within the meta-training support, an assumption that is hard to verify and may fail on genuinely novel conditions. The study in Banday et al. (2025) introduces lightweight pseudo-inverse update to reduce computational burden but lacks convergence analysis, offering no assurance that the adapted model improves on the nominal one.

Beyond these individual limitations, existing methods primarily address how the model should be updated once new measurements become available. The effectiveness of online identification, however, also depends on whether the closed-loop trajectory sufficiently excites the unknown dynamics. Tracking controllers often generate correlated state-input trajectories that provide poor excitation in some directions, while safety during transient adaptation remains largely unaddressed. This motivates treating model adaptation, data acquisition, and safe deployment jointly.

2.2 Active and Safe Model Learning

Active model learning addresses the selection of control actions that improve system identification during operation. Unlike passive online learning, active methods deliberately shape the trajectory to increase the information content of the collected measurements, introducing a trade-off between task performance and model improvement. Within Koopman learning, Abraham and Murphey (2019) combines Koopman identification with information-theoretic active learning, allowing the controller to seek regions that improve the learned representation. However, the Koopman operator is repeatedly reconstructed from locally collected data, making the update sensitive to the conditioning of recent trajectory windows, while general state, input, and collision-avoidance constraints are not explicitly incorporated. Beyond Koopman-based methods, Soloperto et al. (2020) augments nonlinear MPC with an active-learning objective to balance model improvement against task performance. A related perspective is provided by dual control, where control actions influence both the physical state and future model uncertainty. In Baltussen et al. (2025), a Gaussian-process posterior is conditioned on predicted actions and incorporated into contingency-horizon robust MPC, allowing information gathering while maintaining recursive feasibility.Information-seeking actions, however, can drive the system toward less familiar operating regions, making exploration particularly consequential when the model used to plan these actions is uncertain.

Several approaches therefore pair active learning with explicit safety mechanisms. Lew et al. (2022) combine Bayesian meta-learned dynamics with reachability-based chance constraints to obtain high-probability safety guarantees, while Naveed et al. (2025) execute an informative trajectory only after verifying that it satisfies prescribed safety and performance conditions. Prajapat et al. (2025) use GP confidence sets together with reachable-returnable safe sets to obtain finite-time guarantees for safe exploration. These methods show that informative data can be acquired safely, but their guarantees rest on a calibrated probabilistic model of the dynamics or on a verification step performed separately from the controller, both of which add computational burden that is difficult to sustain at the control rates of fast robotic systems.

Control barrier functions (CBFs) provide a complementary mechanism for enforcing safety directly within feedback optimization. Folkestad et al. (2020) combine CBFs with Koopman prediction, using the lifted linear model to reduce the cost of backup-trajectory evaluation while accounting for model error. However, this formulation assumes a fixed learned model. Under continual adaptation, the model itself changes during deployment, so the safety mechanism must account for uncertainty that evolves together with the online updates.

Taken together, the approaches reviewed above address important aspects of active learning and safe exploration, but do not directly link the convergence of online model adaptation to both the information generated by the controller and the uncertainty that must be accommodated for safety. SACK addresses this gap through a contractive Koopman adaptation law with explicit convergence guarantees, whose rate is governed by the regressor Gramian. This directly motivates a D-optimal objective for informative data acquisition, while the same adaptation analysis provides uncertainty bounds for tightening the CBF constraints. Conformal calibration is further introduced as a less conservative alternative to worst-case margins. In this way, the information collected during control determines how rapidly the model improves, while the evolving model accuracy determines how safely further learning can proceed.

3 Adaptation for Koopman Models

3.1 Koopman Operator Theory Preliminaries

Consider a nonlinear autonomous system

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}), \tag{1}$$

where $\boldsymbol{x} \in \mathcal{X} \subset \mathbb{R}^{n_x}$, $\mathcal{X}$ is forward invariant, and $f : \mathcal{X} \to \mathbb{R}^{n_x}$ is Lipschitz continuous. Let Φ^t denote the flow of (1). Sampling with period Δt yields the discrete-time map $\boldsymbol{x}_{k+1} = S(\boldsymbol{x}_k)$, $S := \Phi^{\Delta t} : \mathcal{X} \to \mathcal{X}$. Koopman operator theory (Koopman 1931) represents the nonlinear state dynamics through the evolution of scalar *observables*

$\varphi : \mathcal{X} \to \mathbb{R}$ in a function space $\mathcal{F}$ that is closed under composition with S. The Koopman operator $\mathcal{K} : \mathcal{F} \to \mathcal{F}$ maps each observable to its one-step evolution under the state transition map S,

$$(\mathcal{K}\varphi)(\boldsymbol{x}_k) := (\varphi \circ S)(\boldsymbol{x}_k) = \varphi(\boldsymbol{x}_{k+1}), \tag{2}$$

where $\circ$ denotes function composition. $\mathcal{K}$ is linear, regardless of whether f is linear: the nonlinearity of (1) is traded for the infinite dimension of $\mathcal{F}$. In continuous time, $\{\mathcal{K}^t\varphi = \varphi \circ \Phi^t\}_{t\geq 0}$ forms a semigroup with generator $\mathcal{L}\varphi = \nabla\varphi \cdot f$ for continuously differentiable φ, and $\mathcal{K} = \mathcal{K}^{\Delta t}$. The linear structure is most transparent through Koopman eigenfunctions. If $\mathcal{K}\varphi_\lambda = \lambda\varphi_\lambda$, then

$$\varphi_\lambda(\boldsymbol{x}_k) = \lambda^k\, \varphi_\lambda(\boldsymbol{x}_0), \tag{3}$$

with $\lambda = e^{\mu\Delta t}$ for the corresponding generator eigenvalue μ. Eigenfunctions are therefore intrinsic coordinates in which the nonlinear flow evolves linearly, and any span of finitely many eigenfunctions is a $\mathcal{K}$-invariant subspace (Mauroy et al. 2020; Brunton et al. 2021). Except in special cases, however, no finite-dimensional $\mathcal{K}$-invariant subspace that also contains the state coordinates exists (Brunton et al. 2016), so any finite-dimensional linear model of (1) is an approximation (Korda and Mezić 2018a).

3.1.0.1 Finite-dimensional approximation. In practice, a finite-dimensional approximation of the Koopman operator is constructed by selecting a dictionary of observables $\{\phi_1, \ldots, \phi_p\} \subset \mathcal{F}$, with $\phi_j : \mathcal{X} \to \mathbb{R}$, collected into the *lifting map*,

$$\boldsymbol{z}_k := \psi(\boldsymbol{x}_k) = \big[\phi_1(\boldsymbol{x}_k), \ldots, \phi_p(\boldsymbol{x}_k)\big]^\top \in \mathbb{R}^p. \tag{4}$$

The Koopman operator is then approximated on the span $\mathcal{F}_p := \operatorname{span}\{\phi_1, \ldots, \phi_p\}$ of the selected observables, yielding a finite-dimensional representation of the nonlinear dynamics in the lifted coordinates $\boldsymbol{z}_k$. In general, $\mathcal{F}_p$ is not invariant under $\mathcal{K}$, and therefore the restriction $\mathcal{K}|_{\mathcal{F}_p}$ does not map $\mathcal{F}_p$ into itself. A finite-dimensional operator is therefore obtained by projecting the action of $\mathcal{K}$ back onto $\mathcal{F}_p$: $\mathcal{K}_p := \Pi_p\mathcal{K}|_{\mathcal{F}_p}$, where Π_p denotes the $L^2(\mu)$-orthogonal projection onto $\mathcal{F}_p$ with respect to a sampling measure μ. In lifted coordinates, this compression is represented by

$$\boldsymbol{A}_\mu := \arg\min_{\boldsymbol{A}\in\mathbb{R}^{p\times p}} \mathbb{E}_{\boldsymbol{x}\sim\mu}\big\|\psi(S(\boldsymbol{x})) - \boldsymbol{A}\psi(\boldsymbol{x})\big\|_2^2, \tag{5}$$

whose residual $\psi(S(\boldsymbol{x})) - \boldsymbol{A}_\mu\psi(\boldsymbol{x})$ vanishes μ-almost everywhere if and only if $\mathcal{F}_p$ is $\mathcal{K}$-invariant. Extended Dynamic Mode Decomposition (EDMD) (Li et al. 2017) replaces the expectation by samples, and the empirical solution converges to $\boldsymbol{A}_\mu$ as the number of samples grows (Korda and Mezić 2018b). Two consequences are central to this work. First, $\boldsymbol{A}_\mu$ depends on μ whenever $\mathcal{F}_p$ is not invariant, so the best linear model on a fixed dictionary changes with the region of state space being visited, even if the physics do not. Second, the model is *linear in its parameters*, so estimating it is a linear least-squares problem. Section 3.3 exploits this structure recursively.

3.1.0.2 Controlled systems. Consider the control-affine system

$$\dot{\boldsymbol{x}} = f_0(\boldsymbol{x}) + \sum_{i=1}^{n_u} f_i(\boldsymbol{x})\, u_i, \tag{6}$$

where $\boldsymbol{u} = [u_1, \ldots, u_{n_u}]^\top \in \mathcal{U} \subset \mathbb{R}^{n_u}$, $f_0 : \mathcal{X} \to \mathbb{R}^{n_x}$ is the drift vector field, and $f_i : \mathcal{X} \to \mathbb{R}^{n_x}$ are the control vector fields, all Lipschitz continuous on $\mathcal{X}$. For continuously differentiable ψ, the lifted state evolves as $\dot{\boldsymbol{z}} = \nabla\psi(\boldsymbol{x})f_0(\boldsymbol{x}) + \sum_i \nabla\psi(\boldsymbol{x})f_i(\boldsymbol{x})u_i$. If $\mathcal{F}_p$ is invariant under the Lie derivatives along all f_i, i.e., $\nabla\psi f_i = \boldsymbol{\Lambda}_i\psi$, these dynamics are exactly *bilinear*, $\dot{\boldsymbol{z}} = \boldsymbol{\Lambda}_0\boldsymbol{z} + \sum_i u_i\boldsymbol{\Lambda}_i\boldsymbol{z}$ (Goswami and Paley 2017). The linear predictor used throughout this paper corresponds to the special case in which the input acts on the lifted coordinates through a state-independent gain.

Assumption 1. There exist $\boldsymbol{\Lambda} \in \mathbb{R}^{p\times p}$ and $\boldsymbol{b}_i \in \mathbb{R}^p$, $i = 1, \ldots, n_u$, such that $\nabla\psi(\boldsymbol{x})f_0(\boldsymbol{x}) = \boldsymbol{\Lambda}\psi(\boldsymbol{x})$ and $\nabla\psi(\boldsymbol{x})f_i(\boldsymbol{x}) = \boldsymbol{b}_i$ for all $\boldsymbol{x} \in \mathcal{X}$.

Under Assumption 1 with $\boldsymbol{B}_c := [\boldsymbol{b}_1, \ldots, \boldsymbol{b}_{n_u}]$, the lifted dynamics are linear time-invariant,

$$\dot{\boldsymbol{z}} = \boldsymbol{\Lambda}\boldsymbol{z} + \boldsymbol{B}_c\boldsymbol{u}, \tag{7}$$

and zero-order-hold discretization over Δt yields

$$\hat{\boldsymbol{z}}_{k+1|k} = \boldsymbol{A}\boldsymbol{z}_k + \boldsymbol{B}\boldsymbol{u}_k, \qquad \hat{\boldsymbol{x}}_{k+1|k} = \boldsymbol{C}\hat{\boldsymbol{z}}_{k+1|k}, \tag{8}$$

with $\boldsymbol{A} = e^{\boldsymbol{\Lambda}\Delta t} \in \mathbb{R}^{p\times p}$ and $\boldsymbol{B} = \int_0^{\Delta t} e^{\boldsymbol{\Lambda}s}\,\mathrm{d}s\,\boldsymbol{B}_c \in \mathbb{R}^{p\times n_u}$ the nominal model matrices. The reconstruction matrix $\boldsymbol{C} \in \mathbb{R}^{n_x\times p}$ recovers the state from the lifted state. It is exact whenever the coordinate functions of $\boldsymbol{x}$ lie in $\mathcal{F}_p$, which we enforce by including the state in the dictionary, $\boldsymbol{z} = [\boldsymbol{x}^\top, \phi_{n_x+1}, \ldots, \phi_p]^\top$, so that $\boldsymbol{C} = [\boldsymbol{I}_{n_x}\ \boldsymbol{0}]$ (Appendix 13.1). Here $\hat{\boldsymbol{z}}_{k+1|k}$ is the one-step-ahead prediction produced by the nominal model from the measured lifted state $\boldsymbol{z}_k$ and input $\boldsymbol{u}_k$, and the conditioning "$\cdot\,|k$" denotes a quantity predicted from information available up to step k. The same model arises from the extended-state construction of Korda and Mezić (2018a) and from Dynamic Mode Decomposition with control (Proctor et al. 2016). For robotic systems, Assumption 1 typically holds only approximately, since configuration-dependent input maps, such as $\boldsymbol{M}(\boldsymbol{\theta})^{-1}$ in manipulators or the attitude-dependent thrust direction in quadrotors, make $\nabla\psi f_i$ state dependent. The resulting state-dependent input effect is part of the residual below.

The practical appeal of (8) is twofold. First, predictions over a horizon are affine in the input sequence, so MPC with convex costs and constraints expressed through $\boldsymbol{C}\hat{\boldsymbol{z}}$ reduces to a quadratic program (Korda and Mezić 2018a). Second, with the stacked regressor $\boldsymbol{v}_k := [\boldsymbol{z}_k^\top, \boldsymbol{u}_k^\top]^\top$ and $\boldsymbol{W} := [\boldsymbol{A}, \boldsymbol{B}]$, identification from data pairs $\{(\boldsymbol{v}_j, \boldsymbol{z}_{j+1})\}$ is the linear least-squares problem

$$\min_{\boldsymbol{W}} \sum_j \big\|\boldsymbol{z}_{j+1} - \boldsymbol{W}\boldsymbol{v}_j\big\|_2^2, \tag{9}$$

the controlled analogue of (5). The tuple $\{\psi(\cdot), \boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}\}$ is identified offline from nominal input-output data by jointly learning ψ and minimizing a loss of the form (9). The network architecture is given in Appendix 13.1.

3.1.0.3 Sources of model error Since (8) is built on a finite dictionary and identified from nominal data, the lifted dynamics actually realized satisfy

$$\boldsymbol{z}_{k+1} = \boldsymbol{A}\boldsymbol{z}_k + \boldsymbol{B}\boldsymbol{u}_k + \boldsymbol{\epsilon}_k. \tag{10}$$

To separate the components of $\boldsymbol{\epsilon}_k$, let

$$\boldsymbol{W}^\star = [\boldsymbol{A}^\star, \boldsymbol{B}^\star] := \arg\min_{\boldsymbol{W}} \mathbb{E}\big\|\boldsymbol{z}_{k+1} - \boldsymbol{W}\boldsymbol{v}_k\big\|_2^2 \quad (11)$$

denote the best linear model on the fixed dictionary for the *deployed* dynamics and the region of $\mathcal{X} \times \mathcal{U}$ currently visited. Then

$$\boldsymbol{\epsilon}_k = \underbrace{\boldsymbol{z}_{k+1} - \boldsymbol{W}^\star \boldsymbol{v}_k}_{\text{truncation}} + \underbrace{(\boldsymbol{W}^\star - [\boldsymbol{A}, \boldsymbol{B}])\,\boldsymbol{v}_k}_{\text{parameter mismatch}}. \quad (12)$$

The truncation term is set by the dictionary: it vanishes when $\mathcal{F}_p$ is invariant for the deployed dynamics and Assumption 1 holds exactly, and no update of $(\boldsymbol{A}, \boldsymbol{B})$ can remove it. The mismatch term is linear in the regressor with coefficient $\boldsymbol{W}^\star - [\boldsymbol{A}, \boldsymbol{B}]$. It grows under distributional shift for two reasons: physical changes such as payload or wind alter the deployed dynamics, and, by the μ-dependence noted after (5), operating in a region not covered by the training data changes the best-fitting operator even when the physics are unchanged. Because this component lives in the same lifted coordinates and enters linearly, it can be reduced online by refining $(\boldsymbol{A}, \boldsymbol{B})$ towards $\boldsymbol{W}^\star$, which motivates the adaptive module developed next.

3.2 Adaptive Koopman Module

This section formalizes the perturbation model induced by distributional shift and introduces the corresponding adaptive estimation problem. To model distributional shift, we augment the nominal control-affine system (6) with unknown perturbation terms:

$$\dot{\boldsymbol{x}} = f_0(\boldsymbol{x}) + \tilde{f}_0(\boldsymbol{x}) + \sum_{i=1}^{n_u} \big(f_i(\boldsymbol{x}) + \tilde{f}_i(\boldsymbol{x})\big)u_i, \quad (13)$$

where $\tilde{f}_0, \tilde{f}_i : \mathcal{X} \to \mathbb{R}^{n_x}$ arise from parameter variations, unmodeled dynamics, or environmental disturbances, and are unknown at deployment. In principle such a shift perturbs both the dictionary $\psi(\cdot)$ and the operator matrices $(\boldsymbol{A}, \boldsymbol{B})$, and adapting both would be the ideal response. Re-learning $\psi(\cdot)$ online, however, requires retraining the dictionary at every control step, which is incompatible with real-time operation.

Assumption 2. Fixed lifting map. The perturbed dynamics (13) remain representable within the offline-learned observable subspace $\operatorname{span}\{\phi_1, \ldots, \phi_p\}$. Consequently $\psi(\cdot)$ is held fixed during deployment, and adaptation is performed only over $(\boldsymbol{A}, \boldsymbol{B})$.

Assumption 2 is reasonable for the shifts of practical interest: payload changes, aerodynamic disturbances, and actuator degradation alter how the state evolves rather than which functions of the state are needed to describe it, so the dictionary learned under nominal conditions typically remains adequate while the operator acting on it does not. This is consistent with the empirical findings of Singh et al. (2025a). Restricting adaptation to $(\boldsymbol{A}, \boldsymbol{B})$ additionally preserves a *linear* parameter-estimation problem, admitting the closed-form recursive update and convergence guarantees.

Under Assumption 2, the perturbed system admits the lifted representation

$$\boldsymbol{z}_{k+1} = \boldsymbol{A}_k^* \boldsymbol{z}_k + \boldsymbol{B}_k^* \boldsymbol{u}_k, \qquad \boldsymbol{x}_{k+1} = \boldsymbol{C}\boldsymbol{z}_{k+1}, \quad (14)$$

where $\boldsymbol{W}_k^* := [\boldsymbol{A}_k^*, \boldsymbol{B}_k^*] \in \mathbb{R}^{p \times (p+n_u)}$ denotes the true, generally time-varying, finite-dimensional lifted operator associated with the fixed dictionary $\{\phi_1, \ldots, \phi_p\}$; it is the object the adaptation law estimates. Should the perturbed dynamics leave the learned subspace, so that Assumption 2 fails, the residual $\boldsymbol{\epsilon}_k$ of (10) cannot be removed by operator adaptation alone. In summary, $\psi(\cdot)$ and $\boldsymbol{C}$ are learned offline and held fixed, whereas $(\boldsymbol{A}, \boldsymbol{B})$ are adapted online. Defining the stacked regressor $\boldsymbol{v}_k := [\boldsymbol{z}_k^\top, \boldsymbol{u}_k^\top]^\top \in \mathbb{R}^{p+n_u}$, (14) reads compactly as $\boldsymbol{z}_{k+1} = \boldsymbol{W}_k^* \boldsymbol{v}_k$.

Assumption 3. Bounded operator drift. There exists $\nu \geq 0$ such that $\|\boldsymbol{W}_{k+1}^* - \boldsymbol{W}_k^*\|_F \leq \nu,\ \forall\, k \geq 0$.

Assumption 3 is standard in adaptive estimation (Singh et al. 2025a) and requires only that the true operator vary at a bounded rate. The adaptation law therefore tracks a drifting target, with achievable estimation accuracy fundamentally limited by ν.

Since $\boldsymbol{W}_k^*$ is unknown, we maintain an adaptive estimate $\hat{\boldsymbol{W}}_k = [\hat{\boldsymbol{A}}_k, \hat{\boldsymbol{B}}_k]$, with estimation error $\boldsymbol{E}_k := \boldsymbol{W}_k^* - \hat{\boldsymbol{W}}_k \in \mathbb{R}^{p \times (p+n_u)}$. The resulting one-step prediction error is

$$\boldsymbol{\varepsilon}_{k+1} := \boldsymbol{z}_{k+1} - \hat{\boldsymbol{z}}_{k+1|k} = \boldsymbol{W}_k^* \boldsymbol{v}_k - \hat{\boldsymbol{W}}_k \boldsymbol{v}_k = \boldsymbol{E}_k \boldsymbol{v}_k. \quad (15)$$

The adaptation law below updates $\hat{\boldsymbol{W}}_k$ from closed-loop measurements to drive $\|\boldsymbol{E}_k\|_F$ into a bounded neighbourhood determined by ν.

3.3 Online Adaptation Law

To improve conditioning and noise robustness, the update uses the $w(\geq 1)$ most recent regressor-observation pairs rather than a rank-one correction. Define the stacked matrices

$$\boldsymbol{V}_k := [\boldsymbol{v}_{k-w+1}, \ldots, \boldsymbol{v}_k] \in \mathbb{R}^{(p+n_u) \times w}, \quad (16)$$

$$\boldsymbol{Z}_k^+ := [\boldsymbol{z}_{k-w+2}, \ldots, \boldsymbol{z}_{k+1}] \in \mathbb{R}^{p \times w}, \quad (17)$$

so that each column of $\boldsymbol{Z}_k^+$ is the one-step-ahead observation of the corresponding column of $\boldsymbol{V}_k$. An exponential weighting matrix is introduced to discount the contribution of past observations: $\boldsymbol{\Gamma} := \operatorname{diag}\big(\gamma^{w-1}, \gamma^{w-2}, \ldots, \gamma, 1\big) \in \mathbb{R}^{w \times w}$, where $0 < \gamma \leq 1$ is the forgetting factor. Smaller values of γ place greater emphasis on recent observations, while $\gamma = 1$ corresponds to a uniformly weighted window. The weighted regressor Gramian is defined as $\boldsymbol{G}_k := \boldsymbol{V}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top \in \mathbb{R}^{(p+n_u) \times (p+n_u)}$. By construction, $\boldsymbol{G}_k$ is symmetric positive semidefinite, and becomes positive definite whenever the columns of $\boldsymbol{V}_k$ span $\mathbb{R}^{p+n_u}$, i.e., whenever the persistent excitation condition is satisfied over the current window. The windowed prediction error matrix is defined as $\boldsymbol{E}_k^{\text{pred}} := \boldsymbol{Z}_k^+ - \hat{\boldsymbol{W}}_k \boldsymbol{V}_k \in \mathbb{R}^{p \times w}$.

Assumption 4. There exists $v_{\max} := \sup_{k \geq 0} \|\boldsymbol{v}_k\|_2 < \infty$.

Because the true operator drifts *within* the window (Assumption 3), the columns of $\boldsymbol{Z}_k^+$ are generated by different operators. Indexing the columns of $\boldsymbol{V}_k$ by their sampling times $t_j := k - w + j,\ j = 1, \ldots, w$, the j-th column of $\boldsymbol{Z}_k^+$ satisfies $\boldsymbol{z}_{t_j+1} = \boldsymbol{W}_{t_j}^* \boldsymbol{v}_{t_j}$, and hence the windowed residual admits the exact decomposition

$$\boldsymbol{E}_k^{\text{pred}} = \boldsymbol{E}_k \boldsymbol{V}_k + \boldsymbol{\Xi}_k,\ \boldsymbol{\Xi}_k := \big[\boldsymbol{D}_{k,1}\boldsymbol{v}_{t_1}, \ldots, \boldsymbol{D}_{k,w}\boldsymbol{v}_{t_w}\big], \quad (18)$$

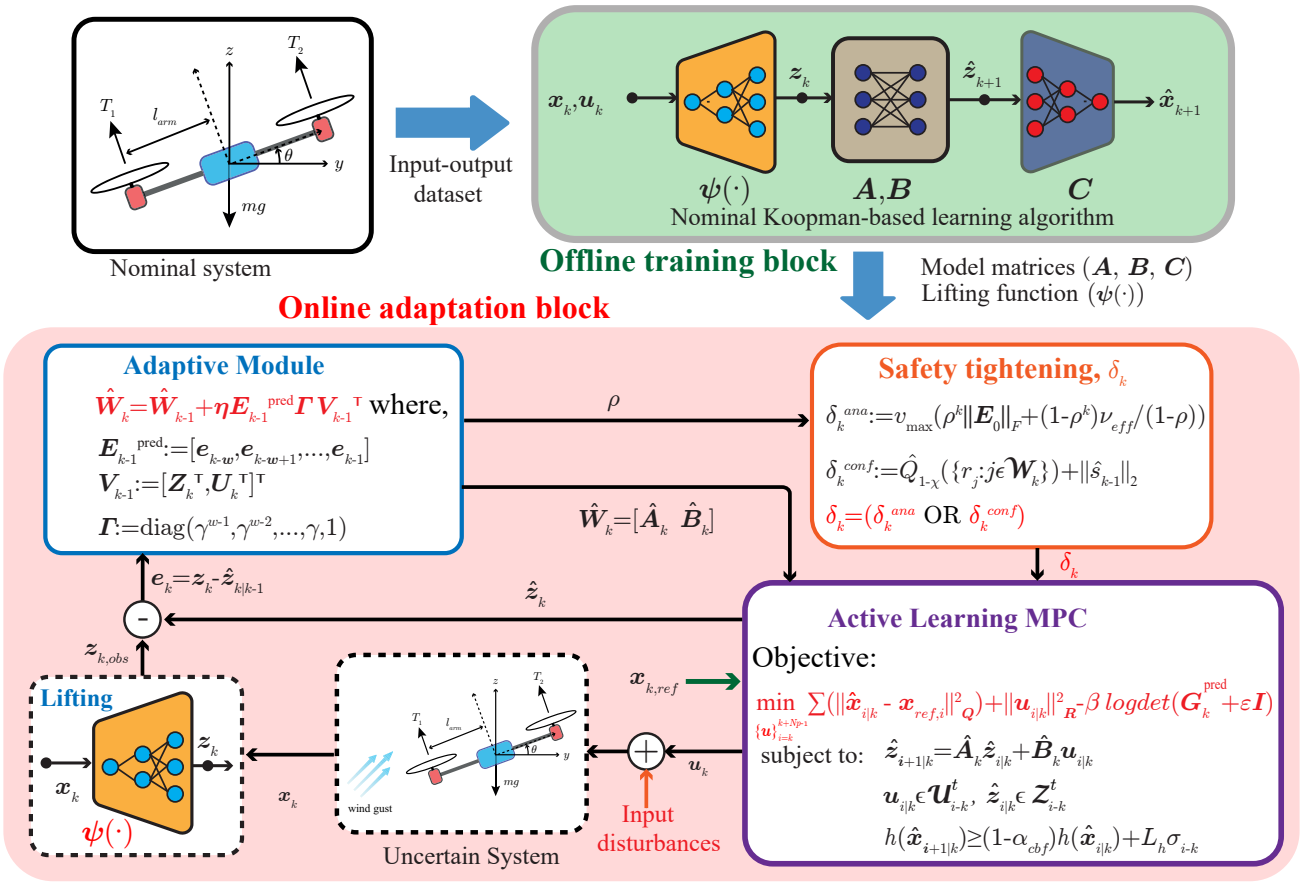


Fig. 1. Overview of the SACK architecture. Offline (top), an autoencoder identifies the lifting map $\psi(\cdot)$ and nominal Koopman matrices $(\boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C})$ from nominal input-output data. Online (bottom), the adaptive module updates $\hat{\boldsymbol{W}}_k = [\hat{\boldsymbol{A}}_k, \hat{\boldsymbol{B}}_k]$ using the contractive adaptation law (19) and windowed prediction errors. The safety-tightening module computes the one-step margin $\delta_{x,k}$ using either the analytical bound (37) or conformal calibration (54). The active-learning MPC (29) jointly performs task tracking, information-seeking excitation, and safety-constrained control under distributional shift

where $\boldsymbol{D}_{k,j} := \boldsymbol{W}^*_{t_j} - \boldsymbol{W}^*_k = -\sum_{i=t_j}^{k-1} \boldsymbol{\Delta}_i$ collects the intra-window operator drift, with $\boldsymbol{\Delta}_i := \boldsymbol{W}^*_{i+1} - \boldsymbol{W}^*_i$, and Assumption 3 gives $\|\boldsymbol{D}_{k,j}\|_F \le (w-j)\,\nu$. For $w = 1$ the drift term vanishes identically, $\boldsymbol{\Xi}_k \equiv \boldsymbol{0}$, and (18) reduces to the exact rank-one identity $\boldsymbol{E}_k^{\text{pred}} = \boldsymbol{E}_k \boldsymbol{v}_k$.

The adaptive Koopman update rule is given by

$$\hat{\boldsymbol{W}}_{k+1} = \hat{\boldsymbol{W}}_k + \eta\, \boldsymbol{E}_k^{\text{pred}} \boldsymbol{\Gamma} \boldsymbol{V}_k^\top, \qquad \eta > 0, \tag{19}$$

where η is a constant step size. Substituting (18), the update (19) is equivalently expressed as

$$\hat{\boldsymbol{W}}_{k+1} = \hat{\boldsymbol{W}}_k + \eta\, \boldsymbol{E}_k \boldsymbol{G}_k + \eta\, \boldsymbol{\Xi}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top, \tag{20}$$

whose first correction term projects the estimation error onto the directions spanned by the current regressor window $\boldsymbol{V}_k$, weighted by $\boldsymbol{\Gamma}$, while the second is a bias induced by intra-window operator drift and vanishes for $w = 1$ or $\nu = 0$.

Proposition 1. *Under (19), the estimation error satisfies*

$$\boldsymbol{E}_{k+1} = \boldsymbol{E}_k(\boldsymbol{I} - \eta \boldsymbol{G}_k) + \tilde{\boldsymbol{\Delta}}_k, \;\; \tilde{\boldsymbol{\Delta}}_k := \boldsymbol{\Delta}_k - \eta\, \boldsymbol{\Xi}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top. \tag{21}$$

Moreover, under Assumptions 3 and 4,

$$\|\tilde{\boldsymbol{\Delta}}_k\|_F \;\le\; \nu_{\text{eff}} := \nu\big(1 + \eta\, v_{\max}^2\, S_\gamma\big), \;\; S_\gamma := \sum_{i=0}^{w-1} i\, \gamma^i. \tag{22}$$

Proof. Substituting (20) into $\boldsymbol{E}_{k+1} = \boldsymbol{W}^*_{k+1} - \hat{\boldsymbol{W}}_{k+1}$ yields $\boldsymbol{E}_{k+1} = (\boldsymbol{W}^*_k - \hat{\boldsymbol{W}}_k) - \eta \boldsymbol{E}_k \boldsymbol{V}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top - \eta \boldsymbol{\Xi}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top + (\boldsymbol{W}^*_{k+1} - \boldsymbol{W}^*_k) = \boldsymbol{E}_k(\boldsymbol{I} - \eta \boldsymbol{G}_k) + \tilde{\boldsymbol{\Delta}}_k$, establishing (21). For (22), write $\boldsymbol{\Xi}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top = \sum_{j=1}^{w} \gamma^{w-j}\, \boldsymbol{D}_{k,j} \boldsymbol{v}_{t_j} \boldsymbol{v}_{t_j}^\top$. Since $\|\boldsymbol{D}\boldsymbol{v}\boldsymbol{v}^\top\|_F = \|\boldsymbol{D}\boldsymbol{v}\|_2 \|\boldsymbol{v}\|_2 \le \|\boldsymbol{D}\|_F \|\boldsymbol{v}\|_2^2$, the triangle inequality together with $\|\boldsymbol{D}_{k,j}\|_F \le (w-j)\nu$ and Assumption 4 gives $\|\boldsymbol{\Xi}_k \boldsymbol{\Gamma} \boldsymbol{V}_k^\top\|_F \le \nu\, v_{\max}^2 \sum_{j=1}^{w} \gamma^{w-j}(w-j) = \nu\, v_{\max}^2\, S_\gamma$, and $\|\tilde{\boldsymbol{\Delta}}_k\|_F \le \|\boldsymbol{\Delta}_k\|_F + \eta\, \nu\, v_{\max}^2 S_\gamma \le \nu_{\text{eff}}$.

The recursion (21) decomposes the error evolution into a linear contraction term governed by the windowed Gramian $\boldsymbol{G}_k$ and an additive perturbation $\tilde{\boldsymbol{\Delta}}_k$ attributable to temporal variation of the true operator, comprising the one-step drift $\boldsymbol{\Delta}_k$ and the intra-window drift bias accumulated across the w samples. When $\boldsymbol{G}_k \succ \boldsymbol{0}$ and the step size satisfies $\rho_k := \|\boldsymbol{I} - \eta \boldsymbol{G}_k\|_2 < 1$, the linear map $\boldsymbol{E}_k \mapsto \boldsymbol{E}_k(\boldsymbol{I} - \eta \boldsymbol{G}_k)$ is strictly contractive. The precise step-size condition ensuring uniform contraction and the resulting exponential convergence bound are established in Lemma 1 and Theorem 1 below.

Lemma 1. Consider (21) with $\boldsymbol{G}_k \succeq 0$, and define the uniform eigenvalue upper bound $\bar{\lambda} := c_\gamma v_{\max}^2 \ge \lambda_{\max}(\boldsymbol{G}_k)\;\forall k$, where $c_\gamma := \sum_{i=0}^{w-1} \gamma^i$, which holds under Assumption 4. Then: (i) if $0 < \eta < 2/\lambda_{\max}(\boldsymbol{G}_k)$, the update is non-expansive, $\|\boldsymbol{E}_{k+1}\|_F \le \|\boldsymbol{E}_k\|_F + \|\tilde{\boldsymbol{\Delta}}_k\|_F$; (ii) if

additionally $\boldsymbol{G}_k \succ 0$, it is strictly contractive, $\|\boldsymbol{E}_{k+1}\|_F \le \rho_k\|\boldsymbol{E}_k\|_F + \|\tilde{\boldsymbol{\Delta}}_k\|_F$ with $\rho_k := \|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2 < 1$; (iii) if $\lambda_{\min}(\boldsymbol{G}_k) \ge \underline{\lambda} > 0$ for all k and $0 < \eta < 2/\bar{\lambda}$, the contraction is uniform: $\|\boldsymbol{E}_{k+1}\|_F \le \rho\|\boldsymbol{E}_k\|_F + \|\tilde{\boldsymbol{\Delta}}_k\|_F$ for all k, with $\rho := \max\{|1-\eta\underline{\lambda}|, |1-\eta\bar{\lambda}|\} < 1$ independent of k.

Proof. By submultiplicativity of the Frobenius norm, $\|\boldsymbol{E}_{k+1}\|_F \le \|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2\|\boldsymbol{E}_k\|_F + \|\tilde{\boldsymbol{\Delta}}_k\|_F$. Since $\boldsymbol{G}_k$ is symmetric positive semidefinite, $\|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2 = \max_i |1 - \eta\lambda_i(\boldsymbol{G}_k)|$. Under $0 < \eta < 2/\lambda_{\max}(\boldsymbol{G}_k)$, $|1 - \eta\lambda_i(\boldsymbol{G}_k)| \le 1$ for all i, proving (i). If $\boldsymbol{G}_k \succ 0$, then $\lambda_i(\boldsymbol{G}_k) > 0$ for all i, hence $|1 - \eta\lambda_i(\boldsymbol{G}_k)| < 1$, establishing (ii). For (iii), note first that $\lambda_{\max}(\boldsymbol{G}_k) \le \|\boldsymbol{V}_k\boldsymbol{\Gamma}\boldsymbol{V}_k^\top\|_2 \le \sum_{j=1}^{w}\gamma^{w-j}\|\boldsymbol{v}_{t_j}\|_2^2 \le c_\gamma v_{\max}^2 = \bar{\lambda}$, so every eigenvalue of $\boldsymbol{G}_k$ lies in $[\underline{\lambda}, \bar{\lambda}]$. Since $\lambda \mapsto |1 - \eta\lambda|$ is convex, it attains its maximum over $[\underline{\lambda}, \bar{\lambda}]$ at an endpoint, giving $\|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2 \le \max\{|1-\eta\underline{\lambda}|, |1-\eta\bar{\lambda}|\} = \rho$. The condition $0 < \eta < 2/\bar{\lambda}$ ensures both endpoint values are strictly less than one, hence $\rho < 1$ independently of k.

Theorem 1. *Suppose (a) Assumptions 3 and 4 hold; (b) $\lambda_{\min}(\boldsymbol{G}_k) \ge \underline{\lambda} > 0$ for all k and the step size satisfies $0 < \eta < 2/\bar{\lambda}$, so that $\|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2 \le \rho < 1$ with ρ as in Lemma 1(iii). Then $\|\boldsymbol{E}_{k+1}\|_F \le \rho\|\boldsymbol{E}_k\|_F + \nu_{\text{eff}}$, and hence*

$$\begin{aligned} \|\boldsymbol{E}_k\|_F &\le \rho^k\|\boldsymbol{E}_0\|_F + \frac{1-\rho^k}{1-\rho}\,\nu_{\text{eff}}, \\ \limsup_{k\to\infty}\|\boldsymbol{E}_k\|_F &\le \frac{\nu_{\text{eff}}}{1-\rho}, \end{aligned} \tag{23}$$

with $\nu_{\text{eff}} = \nu(1 + \eta v_{\max}^2 S_\gamma)$ from (22).

Proof. By Lemma 1(iii) and Proposition 1, $\|\boldsymbol{E}_{k+1}\|_F \le \rho\|\boldsymbol{E}_k\|_F + \nu_{\text{eff}}$. Unrolling the recursion yields

$$\begin{aligned} \|\boldsymbol{E}_k\|_F &\le \rho^k\|\boldsymbol{E}_0\|_F + \sum_{i=0}^{k-1}\rho^i\nu_{\text{eff}} \\ &= \rho^k\|\boldsymbol{E}_0\|_F + \frac{1-\rho^k}{1-\rho}\,\nu_{\text{eff}}. \end{aligned}$$

Taking the limit superior completes the proof.

Remark 1. The drift inflation factor in ν_{eff} is governed by $S_\gamma = \sum_{i=0}^{w-1} i\gamma^i$, the γ-weighted average age of the samples in the window. For $w = 1$, $S_\gamma = 0$ and Theorem 1 recovers the exact rank-one bound with $\nu_{\text{eff}} = \nu$. For $\gamma = 1$, $S_\gamma = w(w-1)/2$ grows quadratically in the window length, whereas for $\gamma < 1$, $S_\gamma \le \gamma/(1-\gamma)^2$ uniformly in w. The forgetting factor therefore actively suppresses the drift-induced bias: larger windows improve the conditioning of $\boldsymbol{G}_k$ (raising $\underline{\lambda}$ and hence the contraction rate) at the cost of an $O(\eta\nu v_{\max}^2 S_\gamma)$ bias, and choosing $\gamma < 1$ caps this cost independently of w. This quantifies the selection of moderate window lengths with exponential forgetting. Under the step-size condition $\eta < 2/\bar{\lambda} = 2/(c_\gamma v_{\max}^2)$, the inflation factor additionally satisfies $\nu_{\text{eff}} \le \nu(1 + 2S_\gamma/c_\gamma)$, where $S_\gamma/c_\gamma \le w - 1$ is the mean sample age.

Theorem 1 provides an explicit bound on the learning-induced model mismatch, which can be interpreted as a bounded additive disturbance and subsequently absorbed by a robust MPC-like controller to ensure robust closed-loop stability. Condition (b) of the theorem requires every window to excite all $p + n_u$ regressor directions, which is impossible when $w < p + n_u$. Section 4 relaxes this condition to excitation over a block of steps, which the active-learning objective is designed to promote.

4 Active Learning Algorithm

The recursion $\boldsymbol{E}_{k+1} = \boldsymbol{E}_k(\boldsymbol{I} - \eta\boldsymbol{G}_k) + \tilde{\boldsymbol{\Delta}}_k$ of Theorem 1 contracts the parameter error only in directions excited within the sliding window. Closed-loop tracking, however, drives the system toward smooth, repetitive trajectories, so the regressors $\boldsymbol{v}_k$ become correlated, yielding an ill-conditioned or low-rank $\boldsymbol{G}_k$, and consequently stalling or reversing online Koopman adaptation precisely when it is needed. This section addresses this data-quality problem by coupling the adaptation law with an active learning strategy that explicitly optimizes the informativeness of the closed-loop trajectory.

4.1 D-Optimal Active Learning Objective

By Lemma 1, for a step size $0 < \eta < 2/\lambda_{\max}(\boldsymbol{G}_k)$ the contraction factor satisfies $\rho_k = \|\boldsymbol{I} - \eta\boldsymbol{G}_k\|_2 = \max_i |1 - \eta\lambda_i(\boldsymbol{G}_k)| < 1$ if and only if $\lambda_{\min}(\boldsymbol{G}_k) > 0$, and the uniform factor of Lemma 1(iii) is $\rho = \max\{|1-\eta\underline{\lambda}|, |1-\eta\bar{\lambda}|\}$. This factor is non-increasing in $\underline{\lambda} = \inf_k \lambda_{\min}(\boldsymbol{G}_k)$, and strictly decreasing while the $\underline{\lambda}$ branch is active, i.e., while $|1-\eta\underline{\lambda}| \ge |1-\eta\bar{\lambda}|$. Once the upper branch dominates, further excitation of the weakest direction no longer improves the guaranteed rate, and ρ is limited by $\bar{\lambda}$ and the step size. Improving the conditioning of $\boldsymbol{G}_k$ is therefore the primary lever on the convergence rate, i.e., when the regressor window is poorly conditioned. This motivates augmenting the task objective with an information-seeking term that targets $\lambda_{\min}(\boldsymbol{G}_k)$ along the closed-loop trajectory.

4.1.0.1 *Excitation over the horizon.*

A single window contains only w regressors, so $\boldsymbol{G}_k$ is singular whenever $w < p + n_u$, and Theorem 1 then gives no contraction. Consecutive windows, however, contain different regressors, so the error still contracts if the windows are informative *together*.

Lemma 2. Let Assumptions 3 and 4 hold and $0 < \eta < 2/\bar{\lambda}$. If, for some $T \in \mathbb{N}$ and $\alpha > 0$,

$$\lambda_{\min}\Big(\sum_{j=k}^{k+T-1}\boldsymbol{G}_j\Big) \ge \alpha \qquad \forall k \ge 0, \tag{24}$$

then $\|\boldsymbol{E}_{k+T}\|_F \le \rho_T\|\boldsymbol{E}_k\|_F + T\nu_{\text{eff}}$ for all k, with a rate $\rho_T < 1$ that depends only on $\eta\alpha$, $\eta\bar{\lambda}$ and T and is given in (61) (Appendix 14.1).

Lemma 2 is Theorem 1 applied to blocks of T steps and for $T = 1$ it reduces to condition (b). Since T windows contain at most $w + T - 1$ distinct regressors, condition (24) needs only $w + T - 1 \ge p + n_u$, so windows shorter than the regressor dimension are admissible. The predictive controller influences exactly the regressors of the next N_p steps, and hence the block $T = N_p$. The information objective is designed to make this block well conditioned.

4.1.0.2 Information objective. A natural choice is the log-determinant of the Gramian of the predicted regressors alone. It ignores the regressors already stored in the window, which the adaptation law keeps using, and it is singular whenever $N_p < p + n_u$. We therefore combine the stored and the planned regressors, each weighted as the adaptation law will weight it. In the window $\boldsymbol{G}_k$ of the next update, the stored regressor $\boldsymbol{v}_{k-m}$, $1 \le m \le w-1$, carries the forgetting weight γ^m, while each planned regressor $\hat{\boldsymbol{v}}_{k+i|k} := [\hat{\boldsymbol{z}}_{k+i|k}^\top, \boldsymbol{u}_{k+i|k}^\top]^\top$ enters its own window $\boldsymbol{G}_{k+i}$ as the newest sample, with weight one. This gives

$$\begin{aligned}\tilde{\boldsymbol{V}}_k(\boldsymbol{U}) &:= \big[\boldsymbol{v}_{k-w+1}, \dots, \boldsymbol{v}_{k-1},\ \hat{\boldsymbol{v}}_{k|k}, \dots, \hat{\boldsymbol{v}}_{k+N_p-1|k}\big],\\ \tilde{\boldsymbol{\Gamma}} &:= \operatorname{diag}\big(\gamma^{w-1}, \dots, \gamma,\ \underbrace{1, \dots, 1}_{N_p}\big),\\ \boldsymbol{S}_k(\boldsymbol{U}) &:= \tilde{\boldsymbol{V}}_k(\boldsymbol{U})\, \tilde{\boldsymbol{\Gamma}}\, \tilde{\boldsymbol{V}}_k(\boldsymbol{U})^\top, \end{aligned} \tag{25}$$

and the information metric

$$J_{\text{info}}(\boldsymbol{U}) := \log\det\big(\boldsymbol{S}_k(\boldsymbol{U}) + \varepsilon \boldsymbol{I}\big), \qquad \varepsilon > 0, \tag{26}$$

where $\varepsilon \boldsymbol{I}$ keeps the matrix positive definite while the window is still being filled. The weights in $\tilde{\boldsymbol{\Gamma}}$ follow from how the adaptation law uses each regressor: the stored regressor $\boldsymbol{v}_{k-m}$ enters the next window $\boldsymbol{G}_k$ with its forgetting weight γ^m, whereas each planned regressor $\hat{\boldsymbol{v}}_{k+i|k}$ enters its own window $\boldsymbol{G}_{k+i}$ as the newest sample, with weight one. Past data are therefore discounted once, by the forgetting factor already present in (19), and planned data are not discounted. The stored columns are measured and hence constant in the optimization; at start-up, columns for times before $k = 0$ are omitted. Because $\tilde{\boldsymbol{V}}_k$ stacks $w-1$ stored and N_p planned regressors, $\boldsymbol{S}_k$ can be nonsingular whenever $w + N_p - 1 \ge p + n_u$, even if neither w nor N_p exceeds $p + n_u$. This is the same necessary condition as for block excitation in Lemma 2 with $T = N_p$. The information term is combined with the tracking objective through the exploration weight $\beta > 0$:

$$J = J_{\text{task}} - \beta\, J_{\text{info}}, \ \text{where} \tag{27}$$

$$J_{\text{task}} := \sum_{i=k}^{k+N_p-1} \|\boldsymbol{C}\hat{\boldsymbol{z}}_{i|k} - \boldsymbol{x}_{\text{ref},i}\|_{\boldsymbol{Q}}^2 + \|\boldsymbol{u}_{i|k}\|_{\boldsymbol{R}}^2. \tag{28}$$

The choice of log-det is motivated by the identity $\log\det(\boldsymbol{S}_k + \varepsilon\boldsymbol{I}) = \sum_i \log(\lambda_i(\boldsymbol{S}_k) + \varepsilon)$. Since $\log(\cdot + \varepsilon)$ is concave and steepest near zero, the criterion penalizes small eigenvalues most heavily. It thus promotes excitation across *all* eigendirections and, because the stored regressors are already included, rewards planned regressors most for exciting directions that are weak in the stored data. This is preferable to simpler heuristics such as maximizing input energy, which may over-excite already-excited directions while leaving others unimproved. The excitation guaranteed by this objective is quantified in Theorem 2 below.

4.2 Constrained Active-Learning MPC

The constrained optimization problem solved at each step k is

$$\min_{\boldsymbol{U}} \sum_{i=k}^{k+N_p-1} \Big(\big\|\boldsymbol{C}\hat{\boldsymbol{z}}_{i|k} - \boldsymbol{x}_{\text{ref},i}\big\|_{\boldsymbol{Q}}^2 + \big\|\boldsymbol{u}_{i|k}\big\|_{\boldsymbol{R}}^2\Big)$$
$$- \beta\, \log\det\Big(\tilde{\boldsymbol{V}}_k(\boldsymbol{U})\, \tilde{\boldsymbol{\Gamma}}\, \tilde{\boldsymbol{V}}_k(\boldsymbol{U})^\top + \varepsilon \boldsymbol{I}\Big) \tag{29a}$$

$$\text{s.t. } \hat{\boldsymbol{z}}_{i+1|k} = \hat{\boldsymbol{A}}_k\, \hat{\boldsymbol{z}}_{i|k} + \hat{\boldsymbol{B}}_k\, \boldsymbol{u}_{i|k},\ i = k, \dots, k+N_p-1, \tag{29b}$$

$$\hat{\boldsymbol{z}}_{k|k} = \boldsymbol{z}_k, \tag{29c}$$

$$\boldsymbol{u}_{i|k} \in \mathcal{U}_{i-k}^t,\ \hat{\boldsymbol{z}}_{i|k} \in \mathcal{Z}_{i-k}^t,\ i = k, \dots, k+N_p-1, \tag{29d}$$

$$\hat{\boldsymbol{z}}_{k+N_p|k} \in \mathcal{Z}_f, \tag{29e}$$

$$h\big(\boldsymbol{C}\hat{\boldsymbol{z}}_{i+1|k}\big) \ge (1 - \alpha_{\text{cbf}})\, h\big(\boldsymbol{C}\hat{\boldsymbol{z}}_{i|k}\big) + L_h \sigma_{i-k},$$
$$i = k, \dots, k+N_p-1, \tag{29f}$$

where $\boldsymbol{U} := \{\boldsymbol{u}_{i|k}\}_{i=k}^{k+N_p-1}$ is the sequence of planned inputs (the decision variable), $\boldsymbol{Q} \succeq \boldsymbol{0}$ and $\boldsymbol{R} \succ \boldsymbol{0}$ weight the tracking error and control effort, $h(\cdot)$ is the CBF defining the safe set $\mathcal{S} = \{\boldsymbol{x} : h(\boldsymbol{x}) \ge 0\}$, $\alpha_{\text{cbf}} \in (0, 1]$ is its decay rate, and L_h is its Lipschitz constant (Assumption 8). The sets $\mathcal{Z}_j^t$ and $\mathcal{U}_j^t$ are the state and input constraint sets at stage $j = i - k$ of the horizon, shrunk to absorb the accumulated prediction error, $\mathcal{Z}_f$ is the terminal set used to guarantee recursive feasibility, and σ_j is the stage-j CBF safety margin. **All are constructed in Section 5.** The sets are precomputed offline from worst-case constants, so the online problem retains the structure of a nominal MPC; only the one-step bound entering σ_j is updated online. Both $\boldsymbol{V}_k^{\text{pred}}(\boldsymbol{U})$ and $\boldsymbol{S}_k(\boldsymbol{U})$ depend on $\boldsymbol{U}$ through the rolled-out lifted states, which is the source of the nonconvexity in (29a) and the stored regressors do not depend on $\boldsymbol{U}$.

Problem (29) is solved in real time by sequential quadratic programming (SQP). Starting from a warm-started nominal input sequence $\{\boldsymbol{u}_{i|k}^{(0)}\}$, SQP iteratively linearizes the CBF inequalities and the quadratic terminal constraint (29e) around the current rollout, yielding affine constraints. The log-det exploration term is handled through its first-order expansion; since the stored columns are constant and the planned ones carry unit weight, its gradient with respect to the planned regressors is $\nabla_{\boldsymbol{V}_k^{\text{pred}}} J_{\text{info}} = 2(\boldsymbol{S}_k + \varepsilon\boldsymbol{I})^{-1}\boldsymbol{V}_k^{\text{pred}}$ with $\boldsymbol{V}_k^{\text{pred}} := [\hat{\boldsymbol{v}}_{k|k}, \dots, \hat{\boldsymbol{v}}_{k+N_p-1|k}]$, so each subproblem is a standard QP.

The following theorem provides a verifiable excitation certificate: it lower-bounds the minimum eigenvalue of $\boldsymbol{S}_k$ for *any* feasible input sequence whose composite cost does not exceed that of a known informative trajectory.

Theorem 2. *Let $n := p + n_u$, let $\mathcal{F} := \{\boldsymbol{U} \in \mathbb{R}^{n_u N_p} :$ (29b)-(29f) hold$\}$ denote the feasible set, and consider $J(\boldsymbol{U}) = J_{\text{task}}(\boldsymbol{U}) - \beta\, J_{\text{info}}(\boldsymbol{U})$, $\beta > 0$, with J_{info} from (26). Suppose:*

- *(A1) **Bounded regressors:** $\|\boldsymbol{v}_t\|_2 \le v_{\max}$ for the stored regressors and $\|\hat{\boldsymbol{v}}_{i|k}\|_2 \le v_{\max}$ for all $i \in \{k, \dots, k + N_p - 1\}$ and all $\boldsymbol{U} \in \mathcal{F}$.*
- *(A2) **Bounded task cost:** $\exists\ \bar{J}$ such that $0 \le J_{\text{task}}(\boldsymbol{U}) \le \bar{J} < \infty$ on $\mathcal{F}$.*

(A3) ***Informative feasible trajectory:*** $\exists\ \boldsymbol{U}^{\mathrm{good}} \in \mathcal{F}$ *and* $\lambda_0 > 0$ *with* $\lambda_{\min}\big(\boldsymbol{S}_k(\boldsymbol{U}^{\mathrm{good}})\big) \geq \lambda_0$.

(A4) ***Sufficient exploration weight:***

$$\frac{\bar{J}}{\beta} < n\log(\lambda_0+\varepsilon) - (n-1)\log(\Lambda+\varepsilon) - \log\varepsilon, \tag{30}$$

where $\Lambda \coloneqq (c_\gamma - 1 + N_p)\, v_{\max}^2$.

Then every $\boldsymbol{U} \in \mathcal{F}$ *satisfying the cost-comparison test*

$$J(\boldsymbol{U}) \ \leq\ J(\boldsymbol{U}^{\mathrm{good}}) \tag{31}$$

satisfies

$$\lambda_{\min}\big(\boldsymbol{S}_k(\boldsymbol{U})\big) \geq \lambda^* \coloneqq \frac{(\lambda_0+\varepsilon)^n}{(\Lambda+\varepsilon)^{n-1}} e^{-\bar{J}/\beta} - \varepsilon > 0. \tag{32}$$

In particular, every global minimizer $\boldsymbol{U}^*$ *of* J *over* $\mathcal{F}$ *satisfies* (31)*, and hence* (32)*.*

Proof. We write $\lambda_j(\cdot)$ for the j-th eigenvalue of a symmetric matrix ($j = 1, \ldots, n$). By (A1), $\lambda_j(\boldsymbol{S}_k(\boldsymbol{U})) \leq \operatorname{tr} \boldsymbol{S}_k(\boldsymbol{U}) = \sum_{m=1}^{w-1} \gamma^m \|\boldsymbol{v}_{k-m}\|_2^2 + \sum_{i=0}^{N_p-1} \|\hat{\boldsymbol{v}}_{k+i|k}\|_2^2 \leq \Lambda$, with $c_\gamma = \sum_{m=0}^{w-1} \gamma^m$ from Lemma 1, for all j and all $\boldsymbol{U} \in \mathcal{F}$. Hence, for any $\boldsymbol{U}$ with $\lambda_{\min}(\boldsymbol{S}_k(\boldsymbol{U})) \leq \delta$, $\delta \geq 0$, monotonicity of $\log(\cdot + \varepsilon)$ gives

$$J_{\mathrm{info}}(\boldsymbol{U}) \leq \phi(\delta) \coloneqq \log(\delta+\varepsilon) + (n-1)\log(\Lambda+\varepsilon), \tag{33}$$

and therefore $J(\boldsymbol{U}) \geq -\beta\phi(\delta)$ by (A2). In contrast, by (A3), $J_{\mathrm{info}}(\boldsymbol{U}^{\mathrm{good}}) \geq n\log(\lambda_0+\varepsilon) =: C_{\mathrm{good}}$, hence $J(\boldsymbol{U}^{\mathrm{good}}) \leq \bar{J} - \beta C_{\mathrm{good}}$. Any $\boldsymbol{U}$ with $\lambda_{\min}(\boldsymbol{S}_k(\boldsymbol{U})) \leq \delta$ therefore violates (31) whenever $-\beta\phi(\delta) > \bar{J} - \beta C_{\mathrm{good}}$, i.e., whenever

$$\delta < \frac{(\lambda_0+\varepsilon)^n}{(\Lambda+\varepsilon)^{n-1}}\, e^{-\bar{J}/\beta} - \varepsilon = \lambda^*, \tag{34}$$

and (A4) is precisely the requirement $\lambda^* > 0$. Hence every $\boldsymbol{U} \in \mathcal{F}$ satisfying (31) has $\lambda_{\min}(\boldsymbol{S}_k(\boldsymbol{U})) \geq \lambda^*$. The final claim follows since a global minimizer satisfies $J(\boldsymbol{U}^*) \leq J(\boldsymbol{U}^{\mathrm{good}})$.

Theorem 2 establishes that the optimization cannot sacrifice all excitation to improve tracking: a trajectory with $\lambda_{\min}(\boldsymbol{S}_k) < \lambda^*$ incurs an information deficit that, under (A4), outweighs any achievable reduction in task cost. Condition (A4) requires β to be large enough relative to $\bar{J}$; λ^* increases with λ_0 and β and decreases with $\bar{J}$ and Λ, and it is quantitatively conservative for large n, its role being to certify that excitation cannot collapse. The certificate also transfers to the realized windows: Corollary 1 in Appendix 14 shows that $\lambda_{\min}(\sum_{j=k}^{k+N_p-1} \boldsymbol{G}_j) \geq \lambda_{\min}(\boldsymbol{S}_k) - 2v_{\max}\sum_{i\geq 1}\vartheta_i$, where ϑ_i is the deviation between planned and realized regressors ($\vartheta_0 = 0$, since the first input is applied exactly).

The problem solved at time k thus certifies the block $\{k, \ldots, k+N_p-1\}$ in (24), and the receding horizon visits every such block. Hence (24), and with it the contraction of Lemma 2, holds whenever the certified level and the deviation term are uniform in k. Since inputs after step k are re-planned, ϑ_i, $i \geq 1$, is not controlled by the time-k problem alone, and bounding it uniformly remains open.

5 Deterministic Safety Guarantees

This section quantifies the uncertainty introduced by model mismatch and online adaptation, and incorporates it into the active-learning MPC through safety tightening. We first derive a deterministic disturbance bound, followed by safety guarantees for the closed loop implementation.

5.1 Disturbance and Model-Update Bounds

Since the update law (19) modifies $\hat{\boldsymbol{W}}_k$ at every step, the true lifted dynamics satisfy $\boldsymbol{z}_{k+1} = \hat{\boldsymbol{W}}_k \boldsymbol{v}_k + \boldsymbol{d}_k$, where $\boldsymbol{d}_k \coloneqq (\boldsymbol{W}_k^* - \hat{\boldsymbol{W}}_k)\boldsymbol{v}_k = \boldsymbol{E}_k \boldsymbol{v}_k$ is the one-step prediction error incurred at time k under the current estimate $\hat{\boldsymbol{W}}_k$. Throughout this section we assume that the excitation condition of either Theorem 1(b) or Lemma 2 holds. Then

$$\|\boldsymbol{E}_k\|_F \leq \bar{E}_k \coloneqq \rho_T^{m_k} \max\{\|\boldsymbol{E}_0\|_F, L\} + \big(1-\rho_T^{m_k}\big)L + (T-1)\,\nu_{\mathrm{eff}}, \tag{35}$$

where $m_k \coloneqq \lfloor k/T \rfloor$ and $L \coloneqq T\nu_{\mathrm{eff}}/(1-\rho_T)$. Here $T = 1$ and $\rho_1 \coloneqq \rho$ under condition (b) of Theorem 1, and ρ_T is as in Lemma 2 under (24); ν_{eff} is from (22). The bound follows from (23) and (62), replacing $\|\boldsymbol{E}_0\|_F$ by $\max\{\|\boldsymbol{E}_0\|_F, L\}$ and the remainder $r\nu_{\mathrm{eff}}$ by $(T-1)\nu_{\mathrm{eff}}$; both replacements make $\bar{E}_k$ non-increasing in k, which recursive feasibility uses below. Hence, under Assumption 4,

$$\|\boldsymbol{d}_k\|_2 \leq \delta_k^{\mathrm{ana}} \coloneqq v_{\max}\bar{E}_k. \tag{36}$$

The analytical bound δ_k^{ana} captures the model-mismatch disturbance entering the true state transition at step k. In physical coordinates, $\boldsymbol{x}_{k+1} - \hat{\boldsymbol{x}}_{k+1|k} = \boldsymbol{C}\boldsymbol{d}_k$, so

$$\|\boldsymbol{C}\boldsymbol{d}_k\|_2 \leq \delta_{x,k}^{\mathrm{ana}} \coloneqq \|\boldsymbol{C}\|_2\, \delta_k^{\mathrm{ana}}. \tag{37}$$

The CBF constraints of (29) use a one-step margin $\delta_{x,k}$ intended to bound $\|\boldsymbol{C}\boldsymbol{d}_k\|_2$: either the analytical choice $\delta_{x,k} = \delta_{x,k}^{\mathrm{ana}}$ or the conformal choice (54) of Section 6.

Recursive feasibility additionally requires bounding the change in the estimated Koopman operator between consecutive iterations. For a fixed regressor $\boldsymbol{v}_k$, this prediction drift is given by $(\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k)\boldsymbol{v}_k$. By (20), the operator increment comprises the error-projection term $\eta\boldsymbol{E}_k\boldsymbol{G}_k$ and the intra-window drift term $\eta\boldsymbol{\Xi}_k\boldsymbol{\Gamma}\boldsymbol{V}_k^\top$. Using $\|\boldsymbol{G}_k\|_2 \leq \bar{\lambda} = c_\gamma v_{\max}^2$ (Lemma 1), the bound $\|\boldsymbol{\Xi}_k\boldsymbol{\Gamma}\boldsymbol{V}_k^\top\|_F \leq \nu\, v_{\max}^2 S_\gamma$ established in Proposition 1, and submultiplicativity of the Frobenius norm,

$$\|\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k\|_F \leq \eta\left(c_\gamma\bar{E}_k + S_\gamma\,\nu\right) v_{\max}^2 =: \mu_k. \tag{38}$$

Since $\|(\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k)\boldsymbol{v}_k\|_2 \leq \|\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k\|_F\|\boldsymbol{v}_k\|_2$, the drift in the prediction for a fixed regressor is bounded by $\mu_k v_{\max}$.

The bounds δ_k^{ana} and μ_k are deterministic functions of k and of the constants ρ_T, T, ν, $\|\boldsymbol{E}_0\|_F$ and $v_{\max}$. The rate ρ_T is determined by the step size η, the window parameters (w, γ) and the excitation level α in (24), and $v_{\max}$ by the compact constraint sets. In contrast, the drift bound ν and the initial error $\|\boldsymbol{E}_0\|_F$ depend on the unknown true operator. In practice, they are replaced by conservative upper estimates, for example ν from the magnitude of the physical perturbations considered and $\|\boldsymbol{E}_0\|_F$ from the offline

validation error of the nominal model. With these estimates, both bounds can be evaluated online. They are, however, conservative when ρ_T is close to one, which occurs when tracking produces correlated regressors and the excitation level α is small. This motivates the data-driven tightening of Section 6.1, which calibrates the margin from observed residuals instead.

5.2 Recursive Feasibility, Forward Invariance, and Boundedness

This subsection establishes recursive feasibility of the MPC problem (29), forward invariance of the safe set $\mathcal{S}$, and ultimate boundedness of the closed-loop prediction error, using the deterministic bounds δ_k^{ana} and μ_k of Section 5.1. Throughout, we use the worst-case constants

$$\begin{aligned} \bar{E} &:= \bar{E}_0 = \max\{\|\boldsymbol{E}_0\|_F, L\} + (T-1)\nu_{\mathrm{eff}}, \\ \bar{\delta}^{\mathrm{ana}} &:= v_{\max}\bar{E}, \qquad \bar{\mu} := \eta\big(c_\gamma \bar{E} + S_\gamma \nu\big) v_{\max}^2, \end{aligned} \tag{39}$$

which dominate $\bar{E}_k$, δ_k^{ana}, and μ_k for all k because $\bar{E}_k$ is non-increasing.

Assumption 5. There exists a known compact convex set $\mathcal{W}^* \subset \mathbb{R}^{p\times(p+n_u)}$ such that the true time-varying Koopman operator satisfies $\boldsymbol{W}_k^* \in \mathcal{W}^*$ for all $k \geq 0$.

Assumption 5 states that, although the true Koopman operator may evolve over time because of distributional shift, model uncertainty, or slowly varying system parameters, it remains confined to a known bounded uncertainty set. The compactness of $\mathcal{W}^*$ guarantees that all admissible operators are uniformly bounded, while convexity permits the use of a common Lyapunov certificate and robust control arguments over the entire family of operators. Since (35) and (39) yield the uniform estimation error bound $\|\boldsymbol{E}_k\|_F \leq \bar{E}$ for all k, the adaptive estimates satisfy

$$\hat{\boldsymbol{W}}_k \in \mathcal{W} := \mathcal{W}^* \oplus \mathcal{B}_F(\bar{E}) \qquad \forall k \geq 0, \tag{40}$$

where $\mathcal{B}_F(\bar{E})$ denotes the Frobenius-norm ball of radius $\bar{E}$. Consequently, all subsequent stability, recursive-feasibility, and safety guarantees need only be established uniformly over the fixed compact set $\mathcal{W}$, rather than for the unknown trajectory $\{\boldsymbol{W}_k^*\}_{k\geq 0}$.

Assumption 6. There exist $\boldsymbol{K}_{\mathrm{ctrl}} \in \mathbb{R}^{n_u\times p}$, $\boldsymbol{P}$ with $\underline{p}\boldsymbol{I} \preceq \boldsymbol{P} \preceq \bar{p}\boldsymbol{I}$, and $\tilde{\alpha} \in (0,1)$ such that, for every $[\boldsymbol{A}, \boldsymbol{B}] \in \mathcal{W}$,

$$(\boldsymbol{A}+\boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}})^\top \boldsymbol{P}\,(\boldsymbol{A}+\boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}}) \;\preceq\; \tilde{\alpha}^2 \boldsymbol{P}. \tag{41}$$

Since $\boldsymbol{P}$ is a common Lyapunov matrix over $\mathcal{W}$, it certifies decay of the *time-varying* transition products: for any $\{\hat{\boldsymbol{W}}_i\} \subset \mathcal{W}$ and $\boldsymbol{A}_{\mathrm{cl},i} := \hat{\boldsymbol{A}}_i + \hat{\boldsymbol{B}}_i \boldsymbol{K}_{\mathrm{ctrl}}$,

$$\big\|\boldsymbol{A}_{\mathrm{cl},k+j-1}\cdots\boldsymbol{A}_{\mathrm{cl},k}\big\|_2 \leq c\,\tilde{\alpha}^{\,j},\; c{:=}\sqrt{\bar{p}/\underline{p}},\; \forall j,k{\geq}0. \tag{42}$$

Given $(c, \tilde{\alpha})$, define the horizon-uniform margin

$$\bar{\delta}^+ := c\,\bar{\delta}^{\mathrm{ana}} + \frac{c\,\bar{\mu}\,v_{\max}}{1-\tilde{\alpha}}, \qquad \bar{\delta}_x^+ := \|\boldsymbol{C}\|_2\,\bar{\delta}^+, \tag{43}$$

accounting for model mismatch and adaptation-induced prediction drift compounded over the horizon. The tightened sets entering (29) are constructed from these constants below.

Assumption 7. Constraint sets and robust terminal ingredients. With $\mathcal{Z} := \{\boldsymbol{z} : \boldsymbol{C}\boldsymbol{z} \in \mathcal{X}\}$, the sets $\mathcal{Z}$ and $\mathcal{U}$ are compact and convex. There exists a terminal set $\mathcal{Z}_f = \{\boldsymbol{z} : \boldsymbol{z}^\top \boldsymbol{P}\boldsymbol{z} \leq \tau\} \subseteq \mathcal{Z}$ such that, for every $[\boldsymbol{A}, \boldsymbol{B}] \in \mathcal{W}$ and $\boldsymbol{z} \in \mathcal{Z}_f$: (i) $\boldsymbol{C}\boldsymbol{z} \in \mathcal{X}$, $\boldsymbol{K}_{\mathrm{ctrl}}\boldsymbol{z} \in \mathcal{U}$; and (ii) $(\boldsymbol{A}+\boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}})(\boldsymbol{z}+\boldsymbol{\zeta}) \in \mathcal{Z}_f$ for every $\|\boldsymbol{\zeta}\|_2 \leq \bar{\delta}^+$.

Both conditions are offline-verifiable: (i) is an ellipsoid-in-polytope containment, and, since (41) gives $\|(\boldsymbol{A}+\boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}})(\boldsymbol{z}+\boldsymbol{\zeta})\|_{\boldsymbol{P}} \leq \tilde{\alpha}\big(\sqrt{\tau}+\sqrt{\bar{p}}\,\bar{\delta}^+\big)$, (ii) holds whenever $\tau \geq \tilde{\alpha}^2\bar{p}\,(\bar{\delta}^+)^2/(1-\tilde{\alpha})^2$. The terminal cost decrease $V_f = \boldsymbol{z}^\top\boldsymbol{P}\boldsymbol{z}$, $V_f((\boldsymbol{A}+\boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}})\boldsymbol{z}) \leq \tilde{\alpha}^2 V_f(\boldsymbol{z})$, follows from (41) and need not be assumed. The candidate terminal input in the feasibility proof is $\kappa_f(\boldsymbol{z}) := \boldsymbol{K}_{\mathrm{ctrl}}\boldsymbol{z}$; the MPC does not impose this law online, and $\boldsymbol{U}$ remains a free decision variable.

Assumption 8. There exists $L_h > 0$ such that $h(\boldsymbol{x}+\boldsymbol{e}) \geq h(\boldsymbol{x}) - L_h\|\boldsymbol{e}\|_2$ for all $\boldsymbol{x}, \boldsymbol{e} \in \mathbb{R}^{n_x}$.

Assumption 8 holds globally for smooth h with bounded gradient, and locally for the circular barrier $h(\boldsymbol{x}) = \|\boldsymbol{x}-\boldsymbol{c}\|_2^2 - d_{\mathrm{safe}}^2$ via the local constant (49). By construction, the MPC (29) enforces the tightened barrier constraint

$$h\big(\hat{\boldsymbol{x}}_{k+1|k}\big) \;\geq\; (1-\alpha_{\mathrm{cbf}})\,h(\boldsymbol{x}_k) + L_h\,\delta_{x,k}, \tag{44}$$

for all $k \geq 0$, with $\alpha_{\mathrm{cbf}} \in (0,1]$ and the one-step margin $\delta_{x,k}$ of Section 5.1. Theorem 3 applies for $\delta_{x,k} = \delta_{x,k}^{\mathrm{ana}}$; the conformal choice is covered by Theorem 4.

Assumption 9. The MPC problem (29), posed with the tightened sets (46), is feasible at $k=0$; $\boldsymbol{x}_0 \in \mathcal{S} := \{\boldsymbol{x} \in \mathcal{X} : h(\boldsymbol{x}) \geq 0\}$ and $\boldsymbol{z}_0 = \psi(\boldsymbol{x}_0) \in \mathcal{Z}$.

Since the predicted trajectories of (29) are confined to $\mathcal{Z}\times\mathcal{U}$, we henceforth take $v_{\max} := \max\{\|[\boldsymbol{z}^\top, \boldsymbol{u}^\top]^\top\|_2 : \boldsymbol{z} \in \mathcal{Z},\ \boldsymbol{u} \in \mathcal{U}\}$, which is finite by compactness (Assumption 7) and bounds realized and predicted regressors alike; in particular, Assumption 4 holds with this constant, and it is the bound used in (A1) of Theorem 2.

5.2.0.1 Tightened constraints. Re-solving the MPC at $k+1$ introduces two perturbations relative to the plan made at k: the realized disturbance $\boldsymbol{d}_k$ and the model update $\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k$. Propagated through the ancillary feedback $\boldsymbol{K}_{\mathrm{ctrl}}$, their effect m steps into the horizon is bounded by

$$b_m := c\,\tilde{\alpha}^{\,m}\bar{\delta}^{\mathrm{ana}} + \frac{c\,\bar{\mu}\,v_{\max}\big(1-\tilde{\alpha}^{\,m}\big)}{1-\tilde{\alpha}} \tag{45}$$

for $m = 0, \ldots, N_p-1$. The disturbance term decays along the horizon at rate $\tilde{\alpha}$, and the update term saturates at $c\bar{\mu}v_{\max}/(1-\tilde{\alpha})$. Accumulating these bounds, $\varpi_j := \sum_{\ell=0}^{j-1} b_\ell$ with $\varpi_0 := 0$, the stage-j constraint sets of (29) are

$$\begin{aligned} \mathcal{Z}_j^{\,t} &:= \mathcal{Z} \ominus \mathcal{B}(\varpi_j), \\ \mathcal{U}_j^{\,t} &:= \mathcal{U} \ominus \mathcal{B}\big(\|\boldsymbol{K}_{\mathrm{ctrl}}\|_2\,\varpi_j\big), \end{aligned} \tag{46}$$

where $\ominus$ is the Pontryagin difference. All quantities are precomputed offline from worst-case constants. Appendix 15 shows that the shifted previous plan, corrected by $\boldsymbol{K}_{\mathrm{ctrl}}$, satisfies these constraints at the next step (Lemma 3). This standard robust-MPC argument (**?**) is adapted here to the time-varying model.

To propagate the barrier constraints along the shifted candidate, the horizon CBF rows are enforced with stage-dependent margins

$$\sigma_i := \delta_{x,k} + \|\boldsymbol{C}\|_2\big(\varpi_{i+1} + (1-\alpha_{\mathrm{cbf}})\varpi_i\big) \tag{47}$$

for $i = 0, \dots, N_p - 1$, i.e., (29) enforces $h(\hat{\boldsymbol{x}}_{k+i+1|k}) \geq (1-\alpha_{\mathrm{cbf}})h(\hat{\boldsymbol{x}}_{k+i|k}) + L_h\sigma_i$ with $\hat{\boldsymbol{x}}_{k|k} = \boldsymbol{x}_k$. Since $\varpi_0 = 0$, the stage-0 row implies the single-step constraint (44), which it supersedes.

Theorem 3. *Suppose Assumptions 5-9 hold, the excitation condition of Theorem 1(b) or of Lemma 2 holds, so that (35) applies, and additionally $\mathcal{Z}_f \oplus \mathcal{B}(\bar{\delta}^+) \subseteq \mathcal{Z}^t_{N_p-1}$, $\boldsymbol{K}_{\mathrm{ctrl}}\big(\mathcal{Z}_f \oplus \mathcal{B}(\bar{\delta}^+)\big) \subseteq \mathcal{U}^t_{N_p-1}$, and, for every $[\boldsymbol{A}, \boldsymbol{B}] \in \mathcal{W}$ and $\boldsymbol{z} \in \mathcal{Z}_f \oplus \mathcal{B}(\bar{\delta}^+)$,*

$$h\big(\boldsymbol{C}(\boldsymbol{A} + \boldsymbol{B}\boldsymbol{K}_{\mathrm{ctrl}})\boldsymbol{z}\big) \geq (1-\alpha_{\mathrm{cbf}})\, h(\boldsymbol{C}\boldsymbol{z}) + L_h\bar{\sigma}, \tag{48}$$

where $\bar{\sigma} := \|\boldsymbol{C}\|_2\big(\bar{\delta}^{\mathrm{ana}} + \varpi_{N_p} + (1-\alpha_{\mathrm{cbf}})\varpi_{N_p-1}\big)$ is the last stage margin (47) evaluated with the worst-case bound. Consider the MPC problem (29) posed with the tightened sets (46), the terminal constraint $\hat{\boldsymbol{z}}_{k+N_p|k} \in \mathcal{Z}_f$, and the CBF rows (47), with measured re-initialization (29c) and applied input $\boldsymbol{u}_k = \hat{\boldsymbol{u}}^*_{k|k}$. Then:

(i) ***Recursive feasibility:*** *with $\delta_{x,k} = \delta^{\mathrm{ana}}_{x,k}$ in (47), the MPC problem remains feasible at every sampling instant.*

(ii) ***Forward safety:*** *with $\delta_{x,k} = \delta^{\mathrm{ana}}_{x,k}$, $\boldsymbol{x}_k \in \mathcal{S}$ for all $k \geq 0$.*

(iii) ***Bounded prediction error:*** *$\|\boldsymbol{z}_{k+1} - \hat{\boldsymbol{z}}^*_{k+1|k}\|_2 \leq \delta^{\mathrm{ana}}_k$ for all k, hence $\limsup_{k\to\infty} \|\boldsymbol{z}_{k+1} - \hat{\boldsymbol{z}}^*_{k+1|k}\|_2 \leq v_{\max}\big(T\nu_{\mathrm{eff}}/(1-\rho_T) + (T-1)\nu_{\mathrm{eff}}\big)$.*

The proof, in Appendix 15, shows that the shifted plan from time k, corrected by feedback, remains feasible at time $k+1$. The margins (45) absorb exactly the disturbance and the model update incurred in one step. The terminal conditions cover the one stage that the shifted plan does not inherit from the previous solution; in particular, (48) requires the inflated terminal set to lie far enough inside the safe set that the ancillary feedback preserves the barrier condition there, which can be checked offline over the ellipsoid. Safety then follows from the stage-0 CBF row, whose margin covers the realized one-step prediction error. The excitation hypothesis is not guaranteed by the controller; the active-learning objective of Section 4 promotes it. If the true operator is constant ($\nu = 0$), Lemma 1(i) with $0 < \eta < 2/\bar{\lambda}$ gives $\|\boldsymbol{E}_k\|_F \leq \|\boldsymbol{E}_0\|_F$ without any excitation, and Theorem 3 holds with $\bar{E}_k \equiv \|\boldsymbol{E}_0\|_F$.

All margins in part (i) are computed offline from $\bar{\delta}^{\mathrm{ana}}$, $\bar{\mu}$, c and $\tilde{\alpha}$; part (ii) additionally uses the online bound $\delta^{\mathrm{ana}}_k = v_{\max}\bar{E}_k$, evaluated with conservative estimates of ν and $\|\boldsymbol{E}_0\|_F$. The deployed conformal tightening is covered by Theorem 4.

Remark 2. The global constant L_h in Assumption 8 may be replaced by the local bound evaluated at the nominal predicted state $\hat{\boldsymbol{x}}_{k+1|k} = \boldsymbol{C}\hat{\boldsymbol{z}}_{k+1|k}$. For the circular barrier $h(\boldsymbol{x}) = \|\boldsymbol{x} - \boldsymbol{c}\|_2^2 - d^2_{\mathrm{safe}}$,

$$L^{\mathrm{local}}_h = 2\big(\|\hat{\boldsymbol{x}}_{k+1|k} - \boldsymbol{c}\|_2 + \delta^{\mathrm{ana}}_{x,k}\big), \tag{49}$$

with $\delta^{\mathrm{ana}}_{x,k}$ from (37). Substituting L^{local}_h for L_h in the stage-0 row preserves part (ii) whenever the resulting problem is feasible, since that argument requires Assumption 8 only at the realized step; the recursive-feasibility argument of part (i) uses the global constant.

Remark 3. The dynamics constraint (29b) involves the nominal model $\hat{\boldsymbol{A}}_k$, $\hat{\boldsymbol{B}}_k$ only; the disturbance enters the formulation exclusively through $\delta_{x,k}$ in the CBF constraint (29f), which absorbs the combined effect of $\boldsymbol{d}_k$ and the model-update perturbation μ_k (Section 5.1). This separation between nominal prediction and disturbance handling is the defining feature of robust CBF-tightened MPC and underpins Theorem 3.

6 Probabilistic Safety Guarantees

6.1 Conformal Safety Tightening

As noted in Section 5, the analytic margin δ^{ana}_k becomes overly conservative as $\rho_T \to 1$, which is precisely the correlated-regressor regime induced by tracking control. The resulting tightened CBF constraints may then be infeasible in cluttered environments, motivating a less conservative margin calibrated directly from observed data.

The safe set is $\mathcal{S} := \{\boldsymbol{x} \in \mathcal{X} : h(\boldsymbol{x}) \geq 0\}$ with $h : \mathbb{R}^{n_x} \to \mathbb{R}$ continuously differentiable. The true state evolves as $\boldsymbol{x}_{k+1} = \hat{\boldsymbol{x}}_{k+1|k} + \boldsymbol{C}\boldsymbol{d}_k$ with $\hat{\boldsymbol{x}}_{k+1|k} = \boldsymbol{C}(\hat{\boldsymbol{A}}_k\boldsymbol{z}_k + \hat{\boldsymbol{B}}_k\boldsymbol{u}_k)$, so the nominal-true discrepancy is governed by $\boldsymbol{s}_k := \boldsymbol{C}\boldsymbol{d}_k$. We track its slowly varying component by the exponential moving average (EMA),

$$\hat{\boldsymbol{s}}_k = \alpha_{\mathrm{EMA}}\hat{\boldsymbol{s}}_{k-1} + (1-\alpha_{\mathrm{EMA}})\boldsymbol{s}_k,\ \alpha_{\mathrm{EMA}} \in (0,1), \tag{50}$$

and define the nonconformity score as the one-step-ahead residual

$$r_k := \big\|\boldsymbol{s}_k - \hat{\boldsymbol{s}}_{k-1}\big\|_2. \tag{51}$$

Since $\hat{\boldsymbol{s}}_{k-1}$ is available before $\boldsymbol{s}_k$ is observed, (51) is a *predictive* score: it quantifies the component of the disturbance not anticipated by the EMA, and is therefore admissible for conformal calibration.

Let $\mathcal{W}_k := \{k - n_{\mathrm{conf}}, \dots, k-1\}$ be a sliding calibration window of the n_{conf} most recent scores, chosen with $n_{\mathrm{conf}} \geq 1/\chi - 1$ for a risk level $\chi \in (0,1)$. For a set $\mathcal{R}$ of n_{conf} scores, $\hat{Q}_{1-\chi}(\mathcal{R})$ denotes its $\lceil (n_{\mathrm{conf}}+1)(1-\chi)\rceil$-th smallest element, the finite-sample corrected quantile of split conformal prediction; the condition on n_{conf} ensures that this index does not exceed n_{conf}. The conformal tightening scalar and its warm-up counterpart are

$$\delta^{\mathrm{conf}}_k := \hat{Q}_{1-\chi}\big(\{r_j : j \in \mathcal{W}_k\}\big) + \big\|\hat{\boldsymbol{s}}_{k-1}\big\|_2, \tag{52}$$

$$\delta^{\mathrm{warm}}_k := \hat{Q}_{1-\chi}\big(\{\|\boldsymbol{s}_j\|_2 : j \in \mathcal{W}_k\}\big). \tag{53}$$

The additive term in (52) is the current EMA magnitude, which is measured online rather than assumed bounded. It accounts for any persistent disturbance bias, such as that induced by a steady wind, without requiring that bias to be small. During warm-up $k < k_{\mathrm{warm}}$, before the EMA has converged, calibration is performed directly on the raw disturbance norms as in (53), for which $\|\boldsymbol{C}\boldsymbol{d}_k\|_2 = \|\boldsymbol{s}_k\|_2$ requires no bias correction. The conformal one-step margin is

therefore

$$\delta_{x,k} = \begin{cases} \delta_k^{\text{warm}} & k < k_{\text{warm}}, \\ \delta_k^{\text{conf}} & k \geq k_{\text{warm}}, \end{cases} \tag{54}$$

calibrated entirely from observed residuals, with $k_{\text{warm}} \geq n_{\text{conf}}$ so that the calibration window is fully populated before (52) is used. In the conformal implementation, the stage margins of (29f) are set to $\sigma_i = \delta_{x,k}$ for all i, i.e., without the horizon offsets of (47); the guarantee of Theorem 4 concerns only the stage-0 row and does not use these offsets.

6.2 Per-Step and Finite-Horizon Coverage

This subsection establishes safety guarantees using the probabilistic bounds formulated in the previous subsection, providing a complementary counterpart to the deterministic guarantees of Theorem 3.

Assumption 10. The scores $\{r_{k-n_{conf}}, \ldots, r_{k-1}, r_k\}$ are exchangeable, i.e., their joint distribution is invariant under permutation.

Theorem 4. *Let Assumptions 8 and 10 hold for $k \geq k_{warm}$, and let $n_{conf} \geq 1/\chi - 1$. Suppose that at each step $k \geq k_{warm}$ the MPC enforces*

$$h\big(\hat{\boldsymbol{x}}_{k+1|k}\big) \geq (1 - \alpha_{cbf})\, h(\boldsymbol{x}_k) + L_h \delta_k^{conf}, \tag{55}$$

with δ_k^{conf} given by (52). Then:

1. Per-step guarantee: *For each $k \geq k_{warm}$,*

$$P(h(\boldsymbol{x}_{k+1}) \geq (1 - \alpha_{cbf})\, h(\boldsymbol{x}_k)) \geq 1 - \chi. \tag{56}$$

2. Finite-horizon guarantee: *If $h(\boldsymbol{x}_{k_{warm}}) \geq 0$, then over any horizon of length $H < 1/\chi$,*

$$P(h(\boldsymbol{x}_k) \geq 0,\ \forall\, k \in \{k_{warm}{+}1, \ldots, k_{warm}{+}H\}) \geq 1 - H\chi. \tag{57}$$

Part (i) follows from the rank argument of split conformal prediction applied to the calibration window, and part (ii) from a union bound over the H steps. The proof is given in Appendix 15.

Remark 4. Assumption 10 is standard in conformal prediction (Vovk et al. 2005; Angelopoulos and Bates 2023). Although exact exchangeability is generally violated in adaptive closed-loop systems, its approximation is promoted by (i) calibrating over a short receding window after the warm-up phase, (ii) scoring against the one-step-ahead EMA prediction $\hat{\boldsymbol{s}}_{k-1}$, so that the score measures only the unanticipated disturbance component and is insensitive to slowly varying bias, and (iii) the contractive adaptation law (Theorem 1), which renders the residual process approximately stationary after convergence. Consequently, the recent nonconformity scores are treated as locally exchangeable, allowing δ_k^{conf} to approximate the $(1-\chi)$ marginal score quantile and thereby justify the approximate per-step coverage in (56). This assumption is local to the receding calibration window and does not require exchangeability over the entire closed-loop trajectory.

Algorithm 1 Safe Active Continual Koopman Control (SACK)

Require: Dataset $\mathcal{D}_{\text{nom}}$, window w, forgetting factor γ, step size η, horizon N_p, exploration weight β, risk level χ, warm-up length k_{warm}
Ensure: Control input $\boldsymbol{u}_k^*$

Offline Training:
Train autoencoder to obtain $\psi(\cdot)$, $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$ using loss $\mathcal{L}_{\text{nom}}$ (60)
Initialize $\hat{\boldsymbol{W}}_0 \leftarrow [\boldsymbol{A}, \boldsymbol{B}]$

Online Execution:
for $k = 1, 2, 3, \ldots$ **do**
  $\boldsymbol{z}_k \leftarrow \psi(\boldsymbol{x}_k)$
  if $k < w$ **then**
    $\hat{\boldsymbol{W}}_k \leftarrow \hat{\boldsymbol{W}}_0$
  else
    Construct $\boldsymbol{V}_{k-1}$, $\boldsymbol{Z}_{k-1}^{+}$, $\boldsymbol{\Gamma}$ using (16)–(17)
    Compute Gramian $\boldsymbol{G}_{k-1}$ and prediction error $\boldsymbol{E}_{k-1}^{\text{pred}}$
    Update $\hat{\boldsymbol{W}}_k$ via contractive adaptation law (19) (index $k-1$)
  end if
  Collect the stored regressors $\boldsymbol{v}_{k-w+1}, \ldots, \boldsymbol{v}_{k-1}$ for (25)
  if using analytical bound **then**
    $\delta_{x,k} \leftarrow \delta_{x,k}^{\text{ana}}$ using (37); stage margins σ_i using (47)
  else if using conformal bounds **then**
    if $k < k_{\text{warm}}$ **then**
      $\delta_{x,k} \leftarrow \delta_k^{\text{warm}}$ using (53)
    else
      Update EMA $\hat{\boldsymbol{s}}_{k-1}$ with $\boldsymbol{s}_{k-1}$ using (50)
      $\delta_{x,k} \leftarrow \delta_k^{\text{conf}}$ using (52)
    end if
    Set $\sigma_i \leftarrow \delta_{x,k}$ for all stages i
  end if
  Solve D-optimal active-learning MPC problem (29a) via SQP
  Apply $\boldsymbol{u}_k^*$ to system and observe $\boldsymbol{x}_{k+1}$
  if using conformal bounds **then**
    Update calibration window $\mathcal{W}_k$ with residual r_k.
  end if
end for

7 Results and Discussion

We validate the proposed SACK algorithm via simulation on a 3R serial manipulator, a planar quadrotor, and a 7-DoF Franka Research 3 (FR3) in high-fidelity Gazebo together with hardware experiments on TurtleBot3. In all the cases, the offline-trained model is evaluated under post-training distributional shifts induced by parameter variations and external disturbances. The online learning performance of the proposed framework is the dynamic prediction error defined as,

$$e_{\text{dyn},i} = \sqrt{\tfrac{1}{N_p} \textstyle\sum_{j=1}^{N_p} \|\hat{\boldsymbol{x}}_{i+j|i} - \boldsymbol{x}_{i+j}\|^2}, \tag{58}$$

which quantifies the prediction error for the dynamics model over the prediction horizon. The adaptation speed is characterized by settling time t_s which is defined as the time required for e_{dyn} to enter and remain within 10% of steady-state value. The closed loop performance is evalueted via the tracking RMSE error amd constraint satisfaction. All controllers use full-state feedback and are implemented using the ACADOS solver (Verschueren et al. 2021) (Intel® Core™ i7-10700, 16GB RAM, Nvidia® Geforce™ RTX 4070 Ti). We compare against four baselines: nominal Koopman with no adaptation (NK), neural adaptive Koopman (NAK) (Singh et al. 2025a), EDMD-based active learning (AcK) (Abraham and Murphey 2019), and GP-based dual MPC (GMPC) (Baltussen et al. 2025). Details of Offline training, network architectures, and hyperparameters are provided in Appendix 13.

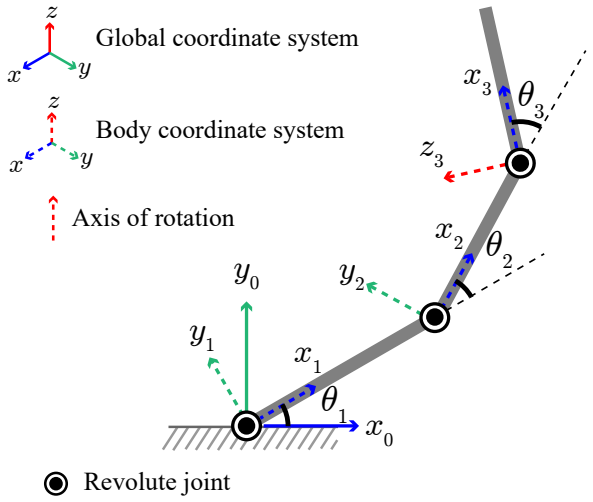


Fig. 2. Schematic diagram illustrating the 3R serial manipulator with a 3D workspace. The axis of rotation for each revolute joint coincides with its respective (body reference) z-axis.

7.1 3R Serial Manipulator

As our first example, we consider a fully actuated 3R manipulator operating a 3D workspace, as illustrated in Figure 2. The coupled dynamics arising from its interconnected links, together with the widespread use of serial manipulators in robotic applications, make this system a suitable benchmark for evaluating the performance of the proposed method

The manipulator follows the standard rigid-body dynamics

$$\boldsymbol{M}(\boldsymbol{\theta})\ddot{\boldsymbol{\theta}} + \boldsymbol{C}(\boldsymbol{\theta}, \dot{\boldsymbol{\theta}})\dot{\boldsymbol{\theta}} + \boldsymbol{G}(\boldsymbol{\theta}) = \boldsymbol{\tau},$$

where $\boldsymbol{M}$, $\boldsymbol{C}$, and $\boldsymbol{G}$ represent the mass, Coriolis, and gravity matrices, respectively. We train the nominal model for link mass $m_i = 0.6$ kg, length $l_i = 1$ m, and inertia $I_i = \text{diag}[0, \frac{m_i l_i^2}{12}, \frac{m_i l_i^2}{12}]$ kgm^2 for $i = 1, 2, 3$. We introduce distribution shifts by modifying the link masses.

7.1.1 Task 1: Safe Exploration in a Constrained Corridor

We first evaluate SACK on constrained trajectory tracking through a narrow corridor containing multiple obstacles. To induce a significant post-training distributional shift, the mass of each manipulator link is reduced by 40% relative to the nominal values used for offline model training. This scenario tests whether the learned model can adapt online while maintaining accurate tracking and satisfying the corridor and obstacle-avoidance constraints. We fix the exploration weight at $\beta = 10^3$, and use conformal tightening to handle transient uncertainty during warmup. Figure 3 shows that SACK traverses the corridor safely without collisions. The executed trajectory deviates from the reference to explore informative regions, consistent with the active learning objective. By improving the conditioning of the adaptation problem, these information-rich trajectories enable the contractive update to rapidly reduce the model error, achieving a settling time of $t_s \approx 0.1$ s (Figure 3c). The conformal margin $\delta_{x,k}$ also decreases rapidly as adaptation progresses (Figure 3d), reflecting the reduction in model mismatch and the corresponding safety margin.

Next, we compare SACK against all the baselines under a more severe 60% mass reduction with $\beta = 10^4$ in an obstacle-free corridor. By design, note that all reference trajectories remain collision-free, so any controller that adapts sufficiently fast should remain inside the corridor without requiring an explicit safety mechanism. As shown in Figure 4, only SACK and GMPC (Baltussen et al. 2025) remain inside the corridor, whereas NK, NAK (Singh et al. 2025a), and AcK (Abraham and Murphey 2019) all violate corridor constraints. SACK adapts with $t_s \approx 0.12$ s, roughly ten times faster than NAK ($t_s \approx 1.17$ s). Moreover, the gradient-based updates of NAK additionally exhibit a secondary error spike at $t = 18.49$ s, which indicates that closed-loop data fails to excite all regressor directions, so corrections don't generalize across the workspace under large distributional shift. AcK Abraham and Murphey (2019) exhibits unsafe behavior because its windowed EDMD re-estimation is sensitive to local data conditioning and carries no convergence guarantee. GMPC (Baltussen et al. 2025) remains safe, but its per-iteration cost is roughly four orders of magnitude higher than SACK's (Figure 4d) on account of the computational burden incurred by nonlinear MPC and online GP updates, precluding real-time deployment. The computational ablation further shows that removing the information objective (SACK-N) reduces runtime for SACK to the order of NK, showing that the law itself introduces negligible computational overhead and the bulk of the computational burden is attributable to the active-learning term. Due to the low cost of the adaption law, SACK, even with the burden of the information active learning objective, achieves computational cost comparable to NAK (Singh et al. 2025a) and AcK (Abraham and Murphey 2019), enabling real-time deployment while providing substantially improved robustness and safety performance under large distribution shift.

7.1.2 Task 2: Tracking Under Parametric Mismatch

To evaluate adaptation independently of obstacle avoidance, all obstacles and corridor constraints are removed. Then, we sweep the system mass from -90% to $+90\%$ while tracking hypotrochoid, petal, and helix reference trajectories. We initialize $\beta = 10^3$ and gradually decay its value as the model converges, thereby shifting the controller from exploration toward tracking precision. GMPC (Baltussen et al. 2025) is

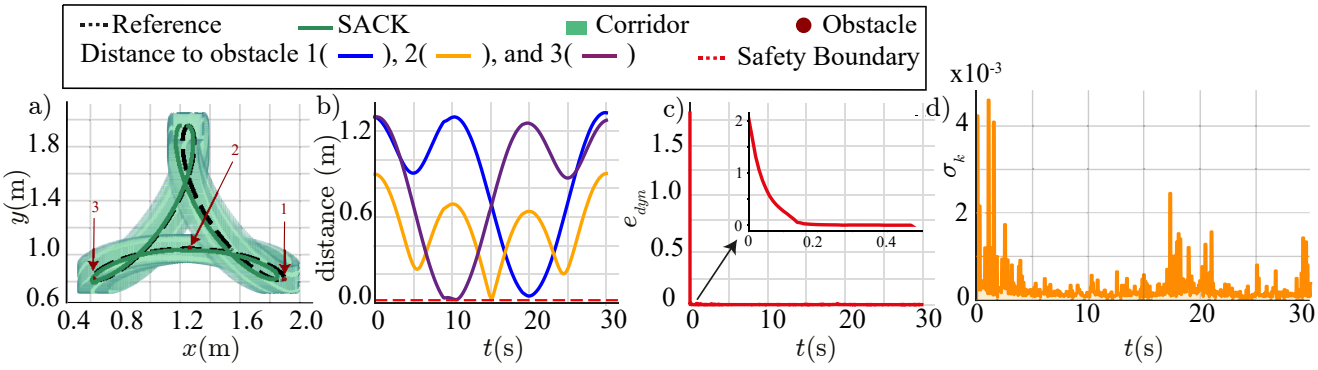


Fig. 3. Constrained tracking performance for SACK for 3R manipulator within narrow corridors with obstacles. The mass of each link is reduced by 40% relative to the nominal training model. a) End-effector trajectory, with red indicating the violation of the corridor boundary. b) Distance to obstacles. c) Evolution of the dynamic prediction error (e_{dyn}). d) Evolution of the conformal one-step margin $\delta_{x,k}$.

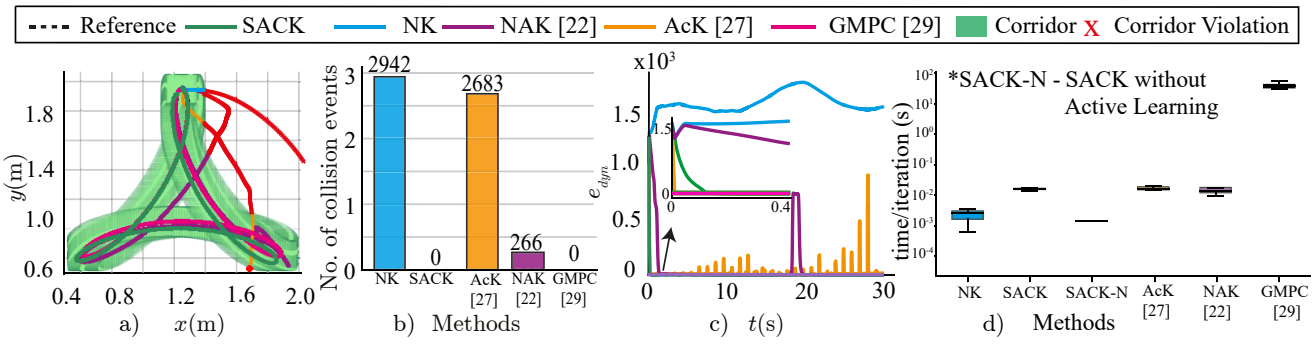


Fig. 4. Comparison of constrained trajectory tracking for the 3R manipulator under distributional shift. The mass of each link is reduced by 60% relative to the nominal model. SACK is compared with NK, AcK (Abraham and Murphey 2019), NAK (Singh et al. 2025a), and GMPC (Baltussen et al. 2025).a) End-effector trajectories within the corridor, with red indicating corridor-boundary violations. b) Number of collision events. c) Dynamic prediction error (e_{dyn}). d) Distribution of computation time per control iteration. SACK-N denotes SACK without active learning ($\beta = 0$).

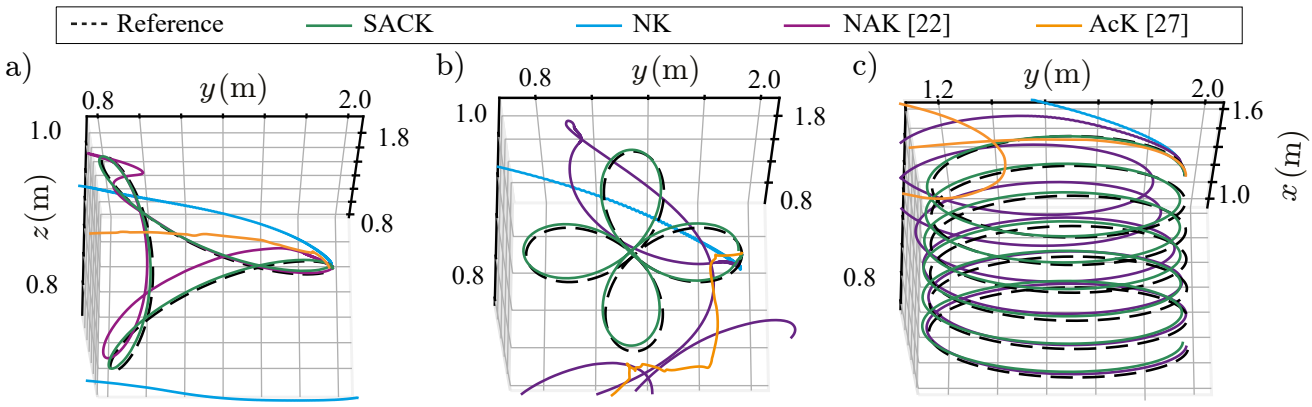


Fig. 5. Trajectory-tracking comparison for the 3R manipulator under a 60% reduction in link masses. SACK is compared with NK, AcK (Abraham and Murphey 2019), and NAK (Singh et al. 2025a) for three reference trajectories: a) Hypotrochoid, b) Petal, c) Helix.

excluded because its high computational cost renders real-time deployment infeasible.

For a representative case of -60% shift (Figure 5), SACK attains RMS tracking errors of 0.0151, 0.0159, and 0.0324 m on the three shapes, respectively, whereas all the baselines show significantly degraded performance. For the full range of distribution shift (Figure 6a), SACK consistently maintains an RMSE of the order of 10^{-2} m, whereas NK exceeds 3 m under moderate-to-large shifts. NAK (Singh et al. 2025a) exhibits inconsistent behaviour with large error spikes that reflect its sensitivity to gradient-update quality. AcK (Abraham and Murphey 2019) shows consistently poor performance.

To isolate the contribution of active data acquisition, we further evaluate SACK without the information objective by setting $\beta = 0$, while retaining the same adaptation law and control formulation (Figure 6b). The resulting SACK-N variant exhibits substantially larger errors under moderate-to-large distributional shifts. This ablation indicates that passive

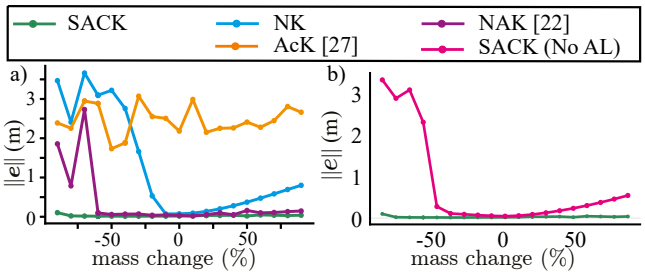


Fig. 6. Tracking RMSE under large variations in the 3R manipulator link masses. Link masses are varied from -90% to $+90\%$ relative to their nominal values. a) Comparison of SACK with NK, AcK (Abraham and Murphey 2019) and NAK (Singh et al. 2025a). b) Abalation of the active-learning objective, comparinig SACK with SACK-N ($\beta = 0$).

closed-loop data do not sufficiently condition the regression problem in these cases, and that actively shaping the collected data plays a critical role in maintaining reliable adaptation under large changes in dynamics.

Taken together, the 3R manipulator studies demonstrate that SACK maintains reliable adaptation across both constrained and unconstrained settings under large post-training variations in system dynamics. The corridor experiments show that rapid model adaptation can be achieved without sacrificing constraint satisfaction, while the mass-sweep study highlights the importance of actively acquiring informative data for maintaining tracking accuracy under severe distributional shifts. Moreover, the computational ablation confirms that the proposed contractive update itself introduces negligible overhead, enabling the combined learning and control framework to remain suitable for real-time implementation.

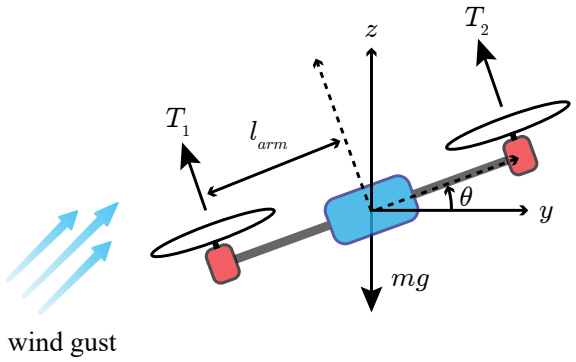


Fig. 7. Schematic diagram illustrating a planar quadrotor. The quadrotor is actuated by two rotors, generating thrusts T_1 and T_2, respectively.

7.2 Planar Quadrotor

Next, we consider a planar quadrotor aerodynamic wind disturbances, as illustrated in Figure 7. The quadrotor provides a qualitatively different benchmark owing to its underactuated and comparatively agile dynamics, allowing us to evaluate whether the proposed framework generalizes beyond serial manipulators. More importantly, the wind disturbance introduces a non parametric model mismatch, in contrast to the fixed parametric variations considered in the preceding experiments. This setting therefore evaluates both the platform-agnostic nature of SACK and its ability to continually adapt when the underlying dynamics evolve during operation.

The planar quadrotor dynamics are given by (Singh et al. 2025a):

$$\begin{bmatrix} \ddot{y} \\ \ddot{z} \\ \ddot{\theta} \end{bmatrix} = \begin{bmatrix} 0 \\ -g \\ 0 \end{bmatrix} + \begin{bmatrix} -\frac{1}{m}\sin\theta & -\frac{1}{m}\sin\theta \\ \frac{1}{m}\cos\theta & \frac{1}{m}\cos\theta \\ -\frac{l_{\text{arm}}}{I} & \frac{l_{\text{arm}}}{I} \end{bmatrix} \begin{bmatrix} T_1 \\ T_2 \end{bmatrix} + \frac{\boldsymbol{F}_w}{m},$$
$$\boldsymbol{F}_w = \begin{bmatrix} Kv_w^2\cos(\alpha_w) & Kv_w^2\sin(\alpha_w) & 0 \end{bmatrix}^\top, \qquad (59)$$

where T_1 and T_2 denote the thrust inputs; m, l_{arm}, I, and g denote the mass, arm length, rotational inertia, and gravitational acceleration, respectively; and $\boldsymbol{F}_w$ represents the wind disturbance with speed v_w, direction α_w, and drag coefficient K. The Koopman model is trained using data collected under nominal operating conditions with no wind ($v_w = 0$). The corresponding system parameters are $m = 2$ kg, $I = 1$ kg m^2, $g = 9.81$ m/s^2, $l_{\text{arm}} = 0.2$ m, and $K = 0.1$ kg/m. Unlike the manipulator experiments, the quadrotor experiences time-varying wind disturbances that induce a drifting perturbation in the true Koopman operator with $\nu > 0$.

7.2.1 Task 1 : Cooperative Navigation Under Model Mismatch We first evaluate SACK on a cooperative navigation task involving two planar quadrotors traversing a shared obstacle field. To introduce a post-training distributional shift, the mass of each quadrotor is increased by 20% and both agents are subjected to a 3 m/s wind disturbance. The exploration weight is set to $\beta = 10$, since the quadrotor dynamics already provide sufficient excitation to keep the regressor Gramian well conditioned without aggressive exploration.

Each quadrotor independently adapts its Koopman model and solves its own MPC, while the safety constraints account for both static obstacles and the neighboring agent. As shown in Figure 8a, both agents successfully traverse the obstacle field without collisions. The inter-agent distance remains above the prescribed safety threshold throughout the maneuver (Figure 8b), demonstrating that continual adaptation does not compromise cooperative safety. At the same time, the dynamic prediction error converges rapidly, with $t_s \approx 0.05$ s (Figure 8c). The low settling time is consistent with the naturally rich excitation generated by the planar quadrotor dynamics allowing the Koopman models to adapt rapidly to the simultaneous mass and wind perturbations.

7.2.2 Task 2: Safe tracking under severe mass and wind disturbances We next consider a more challenging single-quadrotor scenario under a severe distributional shift. The quadrotor is subjected simultaneously to a 30% increase in mass and a 5 m/s wind disturbance while tracking a reference trajectory through an obstacle field.

Figure 9 compares the resulting obstacle-constrained tracking performance. SACK is the only method that maintains constraint satisfaction throughout the task. In contrast, NK, NAK (Singh et al. 2025a), AcK (Abraham and Murphey 2019), and GMPC (Baltussen et al. 2025) all exhibit collisions. The evolution of the dynamic prediction error in

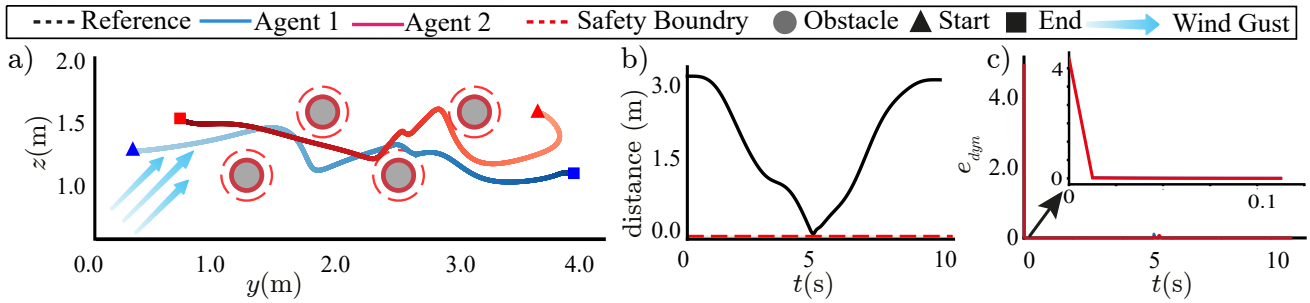


Fig. 8. Cooperative navigation of two planar quadrotors through a shared obstacle field under distributional shift. Each quadrotor is simulated with a 20% mass increase and a $3\,\mathrm{m/s}$ wind disturbance. a) Closed-loop trajectories (arrows indicate the wind direction and dashed circles denote obstacle safety boundaries). b) Inter-agent distance over time. c) Dynamic prediction error e_{dyn}.

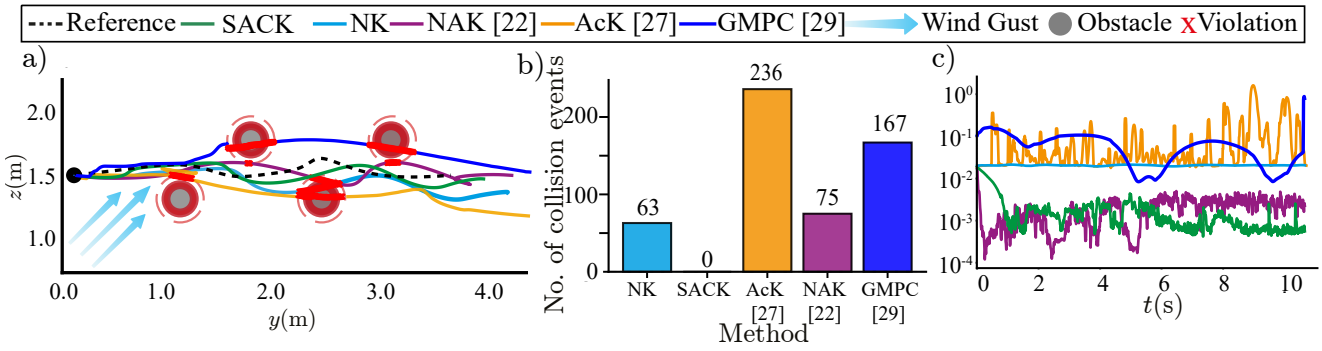


Fig. 9. Comparison of obstacle-constrained trajectory tracking for the planar quadrotor under simultaneous mass and wind disturbances. The quadrotor is evaluated with a 30% mass increase and a $5\,\mathrm{m/s}$ wind disturbance. SACK is compared with NK, AcK (Abraham and Murphey 2019), NAK (Singh et al. 2025a), and GMPC (Baltussen et al. 2025). a) Closed-loop trajectories, with red indicating violations of the obstacle safety boundary. b) Number of collision events. c) Dynamic prediction error e_{dyn}.

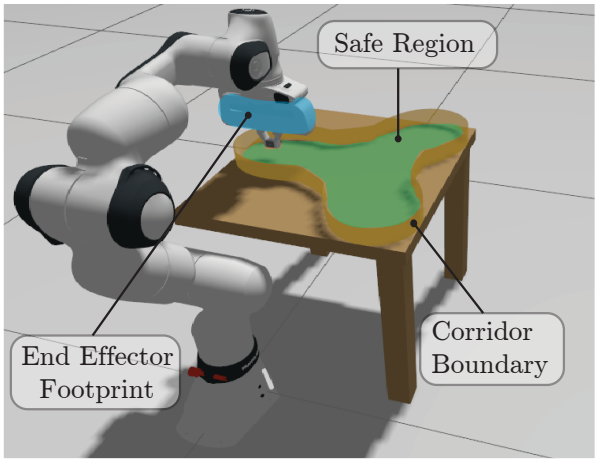


Fig. 10. Simulation setup for corridor-constrained tracking with the Franka Research 3 (FR3). The green region denotes the admissible safe region, while the blue region represents the projected end-effector footprint that must remain entirely within the corridor during tracking.

Figure 9c further explains this behavior. SACK reduces the prediction error rapidly and maintains a low error throughout the maneuver. NAK also reduces the model mismatch, but converges substantially more slowly. This prolonged transient leads to inaccurate predictions during the early part of the maneuver and ultimately to constraint violations. NK exhibits little reduction in model error, while AcK shows large and highly variable prediction-error excursions.

GMPC (Baltussen et al. 2025) exhibits markedly different behavior from that observed in the manipulator experiments. In the manipulator case, the distributional shift is dominated by a fixed change in the link masses, resulting in a comparatively stationary model residual that can be captured effectively by the GP correction model. In the quadrotor case, however, the residual arises from the combined effect of a parametric mass shift and an exogenous wind disturbance that is not explicitly included in the GP regression input. Consequently, similar state-input pairs can correspond to different prediction residuals as the wind disturbance varies, making the resulting mismatch more difficult to represent using a stationary state-input-dependent GP. This is reflected in Figure 9c, where GMPC retains a comparatively large dynamic prediction error over much of the trajectory and fails to adapt sufficiently rapidly to preserve obstacle-constraint satisfaction.

In contrast, SACK updates the Koopman operator directly from the observed lifted prediction errors and does not require the individual sources of model mismatch to be identified separately. The contractive adaptation law rapidly suppresses the prediction error under the combined perturbations, enabling sufficiently accurate closed-loop prediction to preserve obstacle-constraint satisfaction throughout the task. These results highlight the importance of adaptation speed when safety-critical control must be maintained under simultaneous parametric variations and non-prametric external disturbances.

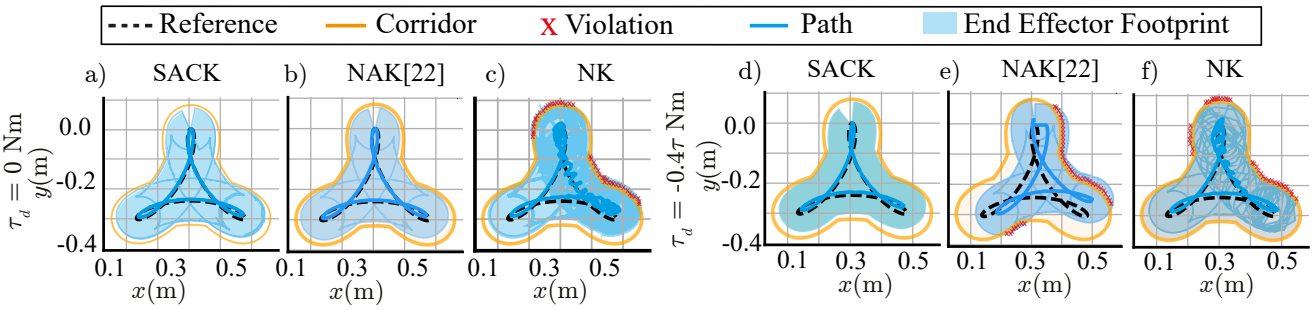


Fig. 11. Corridor-constrained end-effector tracking with the 7-DoF Franka Research 3 (FR3) in Gazebo. The entire end-effector footprint (shaded blue) is required to remain within the admissible corridor. Panels (a)-(c) show the nominal condition, whereas panels (d)-(f) apply the resistive joint-torque disturbance $\boldsymbol{\tau}_d = -0.4\boldsymbol{\tau}$. (a,d) SACK; (b,e) NAK Singh et al. (2025a); (c,f) NK. Red crosses denote corridor violations.

7.3 Franka Research 3 (Gazebo)

Next, we consider a 7-DoF Franka Research 3 (FR3) manipulator in high-fidelity Gazebo simulation. The task is to track a reference trajectory through a narrow three-dimensional corridor while ensuring that the entire projected footprint of the end effector remains inside the prescribed safe region, as illustrated in Figure 10. Unlike the point-wise constraints considered in the previous experiments, this formulation imposes a set-valued geometric constraint. Even if the end-effector reference point lies inside the corridor, the motion is considered unsafe whenever any portion of the end-effector footprint crosses the corridor boundary. This provides a more demanding test of constraint satisfaction in a higher-dimensional robotic system. GMPC (Baltussen et al. 2025) and AcK (Abraham and Murphey 2019) are omitted from this study owing to their high computational cost and poor adaptation performance, respectively, as established in the preceding experiments.

We evaluate the controller under both nominal and perturbed dynamics. In the nominal case, the FR3 operates using the same dynamics distribution represented during offline training. In the perturbed case, a resistive joint-torque disturbance $\boldsymbol{\tau}_d = -0.4\boldsymbol{\tau}$ is applied during execution, creating a substantial post-training model mismatch. The objective is therefore not only to preserve accurate end-effector tracking, but also to ensure that the complete end-effector footprint remains within the admissible region while the learned model adapts to the altered dynamics.

Under nominal conditions, SACK tracks the reference accurately without unnecessary exploratory deviations while maintaining the footprint inside the corridor. Under the resistive torque disturbance, SACK continues to satisfy the geometric safety constraint as the adaptive Koopman model compensates for the altered dynamics and the conformal tightening adjusts to the residual model uncertainty. In contrast, NAK (Singh et al. 2025a) violates the corridor boundary under the perturbed dynamics, while NK fails to maintain constraint satisfaction in both the nominal and disturbed cases.

These results demonstrate that SACK extends to higher-dimensional manipulation tasks with stricter geometric safety requirements while retaining effective adaptation under significant model mismatch.

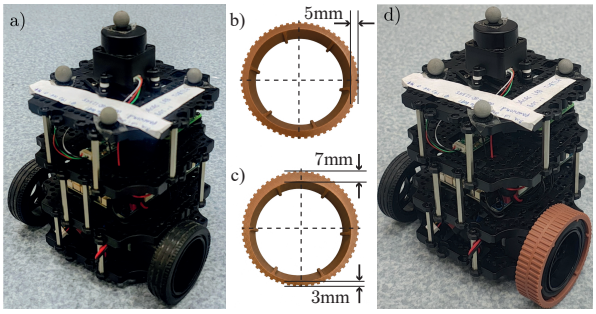


Fig. 12. TurtleBot3 Burger hardware platform and wheel attachments used to induce model mismatch. a) Nominal configuration. b) Axisymmetric wheel ring, which increases the effective radius of the left wheel. c) Eccentric wheel ring, which induces periodic variation in the effective wheel radius and uneven rolling. d) TurtleBot3 with a ring mounted on the left wheel.

Table 1. Hardware tracking and safety performance over 20 trials. Entries report the average tracking RMSE (m) and average number of time steps spent in the unsafe set for the open-circuit (OC) and closed-circuit (CC) tracks.

	SACK	NAK	NK
OC	**0.244, 0**	1.59, 384	1.26, 76
CC	**0.112, 0**	0.984, 389	1.329, 669

7.4 Hardware Experiments: TurtleBot3

We finally deploy SACK on a TurtleBot3 Burger mobile robot (Figure 12), using a Koopman model trained entirely offline in Gazebo. This experiment provides a qualitatively different validation from the preceding simulation studies because the deployment mismatch is no longer prescribed exactly by known parameter variations or disturbance models. Instead, the controller must contend with the inherent sim-to-real gap together with additional unmodeled hardware effects that arise only during physical operation. The TurtleBot3 therefore serves as a test of whether SACK can retain reliable adaptation and safety when the source and structure of the model mismatch are not known a priori.

Three configurations of increasing mismatch are considered: the nominal robot, an axisymmetric ring mounted on one wheel that alters its effective radius and introduces asymmetric kinematic mismatch, and an eccentric ring that produces periodically varying wheel-ground interaction and hence a time-varying disturbance. These perturbations

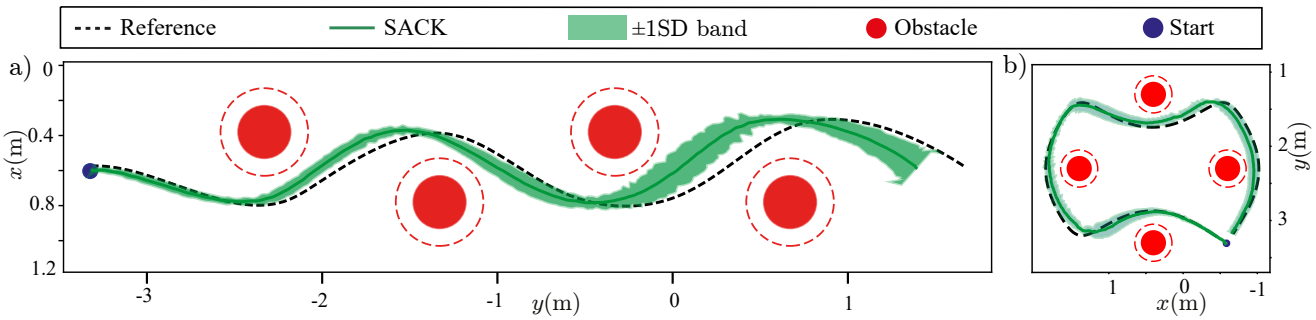


Fig. 13. Distribution of SACK trajectories across 20 TurtleBot3 hardware trials. Trials are pooled across the nominal, axisymmetric-ring, and eccentric-ring configurations. The solid curve denotes the mean trajectory, and the shaded region represents ± 1 standard deviation in the path-normal direction. (a) Open-circuit track. (b) Closed-circuit track.

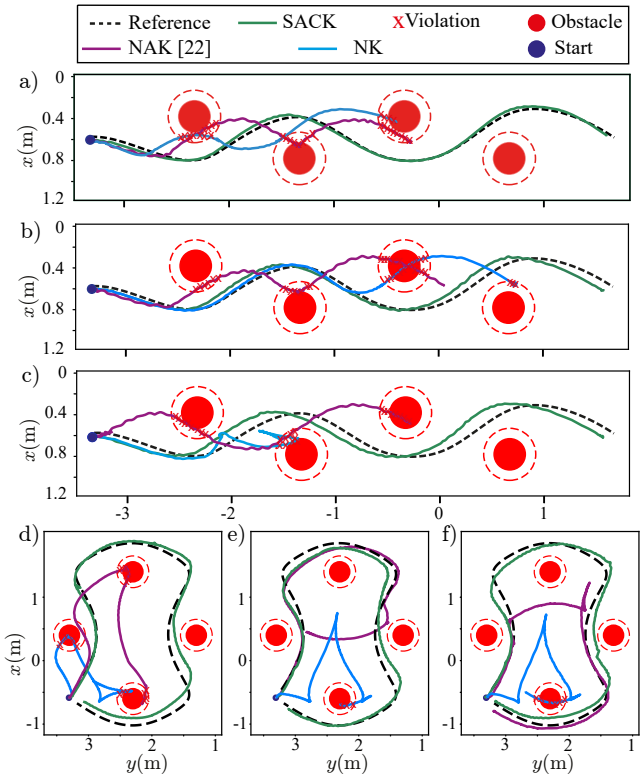


Fig. 14. Representative TurtleBot3 hardware trials under increasing model mismatch. Results are shown for SACK, NAK (Singh et al. 2025a), and NK on the open-circuit track (a-c) and closed-circuit track (d-f). (a,d) Nominal robot. (b,e) Axisymmetric ring mounted on the left wheel. (c,f) Eccentric ring mounted on the left wheel. Red crosses denote violations of the obstacle safety margins.

emulate common deployment effects such as wheel-radius mismatch, uneven wear, actuator imbalance, and irregular rolling behavior. These configurations emulate practical deployment effects such as payload asymmetry, uneven wheel wear, actuator imbalance, and irregular terrain-induced disturbances. The three configurations are evaluated on both open-circuit (OC) and closed-circuit (CC) obstacle courses.

Table 1 and Figure 13 summarize performance over 20 hardware trials across the different configurations and tracks. SACK maintains zero unsafe-set violations while achieving lower tracking RMSE than both baselines, demonstrating that the safety performance is repeatable across trials rather than arising from isolated successful runs. The OC-CC comparison further illustrates how performance improves as online adaptation accumulates data. The CC track is longer, allowing SACK to collect more informative closed-loop samples and progressively refine the Koopman model during deployment, reflected in the reduced RMSE and the tightening trajectory spread visible in the last portion of Figure 13.

Figure 14 provides representative trajectories for SACK and the two baselines under each hardware configuration. Even for the nominal robot, NK and NAK (Singh et al. 2025a) exhibit safety-margin violations due to the sim-to-real mismatch, with performance degrading further under the wheel perturbations. In contrast, SACK completes all evaluated trials without entering the unsafe set, including under the eccentric-wheel configuration where the mismatch varies continuously during motion. These experiments demonstrate that the proposed continual adaptation framework remains effective when the deployment mismatch is unmodeled, hardware-induced, and time varying.

Overall, the simulation and hardware results demonstrate that SACK generalizes across a broad range of robotic systems and post-training distributional shifts, including large parametric variations, nonparametric external disturbances, higher-dimensional geometric constraints, sim-to-real mismatch, and unmodeled hardware effects.

8 Discussion and Future Work

The central tension in continual model adaptation is that the system must explore to learn, yet exploration is most hazardous precisely when the model is least certain. SACK addresses this problem by tightly coupling Koopman adaptation, active data acquisition, and safety-aware control within a unified online framework. The proposed closed-form Koopman update yields error dynamics whose contraction is governed by the regressor Gramian, with exponential convergence under persistent excitation and an ultimate error bound when the underlying Koopman operator drifts over time. This relationship provides the central mechanism of the framework: the information content of the closed-loop data determines how rapidly the model can adapt.

The active-learning component follows directly from this observation rather than being introduced as an auxiliary exploration heuristic. The D-optimal objective shapes future state-input trajectories to improve the conditioning of the regressors used by the adaptive update, and the accompanying analysis establishes conditions under which the resulting

optimization prevents the predicted excitation from collapsing. The 3R manipulator experiments illustrate the practical importance of this coupling. Retaining the same adaptation law while removing the information objective leads to substantially poorer performance under large distributional shifts, showing that an effective update law alone is insufficient when the task trajectory does not provide sufficiently informative data. Thus, continual adaptation depends jointly on how model parameters are updated and on the trajectories from which the update data are obtained.

A second implication concerns the interaction between adaptation and safety. Because the Koopman model evolves during deployment, the uncertainty entering the safety constraints is itself time varying. SACK explicitly connects this uncertainty to the adaptation analysis. The deterministic formulation propagates the derived estimation-error and model-update bounds into tightened MPC and CBF constraints, yielding recursive-feasibility and forward-invariance guarantees under the stated assumptions. These analytical margins can, however, become conservative when excitation is weak or the required worst-case quantities are estimated conservatively, potentially rendering the MPC problem infeasible in tightly constrained environments. The conformal formulation provides a complementary alternative by calibrating the one-step margin directly from recently observed prediction residuals. It therefore trades worst-case robustness for less conservative probabilistic coverage. Importantly, these two formulations should be viewed as complementary operating regimes rather than interchangeable guarantees. The former provides deterministic robustness under explicit boundedness and excitation assumptions, while the latter relies on the exchangeability of the calibration residuals.

The experiments further show that the relevance of rapid adaptation extends beyond fixed parametric uncertainty. SACK is evaluated under large manipulator mass variations, simultaneous mass and wind disturbances for an underactuated planar quadrotor, resistive joint torques on a 7-DoF FR3, and sim-to-real and hardware-induced mismatch on a TurtleBot3. A recurring observation across these settings is that eventual model accuracy alone is not sufficient for safety-critical operation. Slow transient adaptation can produce inaccurate predictions for long enough to violate constraints before the model converges. In contrast, directly adapting the effective lifted dynamics allows SACK to respond to both parametric changes and disturbances whose physical source is not explicitly represented in the model. The hardware results are particularly relevant in this regard because the mismatch is no longer prescribed analytically, but instead arises from the sim-to-real gap and unmodeled physical effects encountered during deployment.

These observations also suggest a broader perspective on offline Koopman learning. Rather than requiring a single Koopman model to remain accurate over every condition that may be encountered after deployment, the offline model can provide a nominal lifted representation that is subsequently refined as the operating conditions change. This shifts part of the modeling burden from global offline generalization to continual local refinement of the dynamics. Such a strategy can be particularly attractive for robotic systems, where exhaustively collecting training data across all payloads, disturbances, environments, and hardware conditions is generally impractical. The present results indicate that substantial deployment mismatch can be compensated by adapting the Koopman operator while retaining the nominal observable coordinates, provided that the perturbed dynamics remain sufficiently representable within the learned lifting space.

This requirement also identifies an important limitation of the present framework. The lifting map $\psi(\cdot)$ is fixed after offline training, and therefore distributional shifts that substantially alter the appropriate observable subspace cannot in general be compensated by adapting (A, B) alone. Jointly adapting the lifting representation and the Koopman operator while retaining computational efficiency and meaningful convergence guarantees is an important direction for future work.

Further extensions arise from the practical safety and optimization components. The deterministic safety formulation relies on worst-case bounds on the model mismatch, operator drift, and adaptation-induced prediction error, which can become conservative when these quantities are only coarsely estimated. The resulting constraint tightening may significantly restrict the feasible region and can render the MPC problem infeasible in tightly constrained environments. Developing tighter online uncertainty estimates and less conservative propagation schemes would therefore improve the practical applicability of the deterministic safety guarantees. The exploration weight β is also selected manually or scheduled empirically; adapting it online according to quantities such as Gramian conditioning, model uncertainty, or task progress could provide a more principled exploration–exploitation balance. Finally, although the closed-form Koopman update itself introduces little computational overhead, the information-seeking MPC is nonconvex and its SQP solution becomes increasingly demanding as the system and information dimensions grow. More scalable information objectives, structured optimization methods, and distributed formulations will therefore be important for extending SACK to substantially higher-dimensional and larger multi-agent robotic systems.

Overall, SACK establishes a direct connection between the information obtained during control, the convergence of the continually adapted Koopman model, and the uncertainty that must be accommodated for safe operation. The simulation and hardware studies indicate that exploiting this connection enables continual model refinement across a broad range of robotic systems and post-training distributional shifts, providing a step toward learning-based controllers that can continue to adapt during deployment without decoupling learning from safety.

9 Authorship Contribution

All three authors contributed to the conception, design, and development of the proposed algorithm. C.K.S. and R.S. implemented, tested, and refined the method. The manuscript was written by R.S and C.K.S with support from J.K.

10 Data availability

The data used in the study can be found at : Data Repository. The code used for generating all the results in the study can be found at : GitHub Repository

11 Statement and Declarations

11.1 Ethical Considerations

Not Applicable

11.2 Consent to Participate

Not Applicable

11.3 Consent for publication

Not Applicable

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

This work is funded in part by the Science and Engineering Research Board (SERB) Core Research Grant.

References

Abraham I and Murphey TD (2019) Active learning of dynamics for data-driven control using koopman operators. *IEEE Transactions on Robotics* 35(5): 1071–1083.

Angelopoulos AN and Bates S (2023) Conformal prediction: A gentle introduction. *Foundations and Trends in Machine Learning* 16(4): 494–591.

Baltussen T, Heemels M and Katriniok A (2025) Dual mpc for active learning of nonparametric uncertainties. *arXiv preprint arXiv:2511.08542* .

Banday B, Sah CK and Keshavan J (2025) Event-based adaptive koopman framework for optic flow-guided landing on moving platforms. *arXiv preprint arXiv:2501.16868* .

Bruder D, Bombara D and Wood RJ (2025) A koopman-based residual modeling approach for the control of a soft robot arm. *The International journal of robotics research* 44(3): 388–406.

Bruder D, Fu X, Gillespie RB, Remy CD and Vasudevan R (2020) Data-driven control of soft robots using koopman operator theory. *IEEE transactions on robotics* 37(3): 948–961.

Brunton SL, Brunton BW, Proctor JL and Kutz JN (2016) Koopman invariant subspaces and finite linear representations of nonlinear dynamical systems for control. *PLoS ONE* 11(2): e0150171.

Brunton SL, Budišić M, Kaiser E and Kutz JN (2021) Modern koopman theory for dynamical systems. *arXiv preprint arXiv:2102.12086* .

Champion K, Lusch B, Kutz JN and Brunton SL (2019) Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences* 116(45): 22445–22451.

Dixon WE, Behal A, Dawson DM and Nagarkatti SP (2003) Nonlinear control of engineering systems. *Boston, MA: Birkäuser* .

Folkestad C, Chen Y, Ames AD and Burdick JW (2020) Data-driven safety-critical control: Synthesizing control barrier functions with koopman operators. *IEEE Control Systems Letters* 5(6): 2012–2017.

Folkestad C, Wei SX and Burdick JW (2022) Koopnet: Joint learning of koopman bilinear models and function dictionaries with application to quadrotor trajectory tracking. In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1344–1350.

Goswami D and Paley DA (2017) Global bilinearization and controllability of control-affine nonlinear systems: A koopman spectral approach. In: *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*. pp. 6107–6112.

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

Korda M and Mezić I (2018a) Linear predictors for nonlinear dynamical systems: Koopman operator meets model predictive control. *Automatica* 93: 149–160.

Korda M and Mezić I (2018b) On convergence of extended dynamic mode decomposition to the Koopman operator. *Journal of Nonlinear Science* 28(2): 687–710.

Krolicki A, Rufino D, Zheng A, Narayanan SS, Erb J and Vaidya U (2022) Modeling quadruped leg dynamics on deformable terrains using data-driven koopman operators. *IFAC-PapersOnLine* 55(37): 420–425.

Lew T, Sharma A, Harrison J, Bylard A and Pavone M (2022) Safe active dynamics learning and control: A sequential exploration–exploitation framework. *IEEE Transactions on Robotics* 38(5): 2888–2907.

Li F, Abuduweili A, Sun Y, Chen R, Zhao W and Liu C (2024) Continual learning and lifting of koopman dynamics for linear control of legged robots. *arXiv preprint arXiv:2411.14321* .

Li Q, Dietrich F, Bollt EM and Kevrekidis IG (2017) Extended dynamic mode decomposition with dictionary learning: A data-driven adaptive spectral decomposition of the koopman operator. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 27(10): 103111.

Lusch B, Kutz JN and Brunton SL (2018) Deep learning for universal linear embeddings of nonlinear dynamics. *Nature communications* 9(1): 4950.

Manaa ZM, Abdallah AM, Abido MA and Ali SSA (2024) Koopman-lqr controller for quadrotor uavs from data. In: *2024 IEEE International Conference on Smart Mobility (SM)*. IEEE, pp. 153–158.

Mauroy A, Susuki Y and Mezic I (2020) *Koopman operator in systems and control*, volume 484. Springer.

Naveed KB, Agrawal DR and Panagou D (2025) A formal gatekeeper framework for safe dual control with active exploration. *arXiv preprint arXiv:2510.06351* .

Ogata K et al. (2010) *Modern control engineering*, volume 5. Prentice hall Upper Saddle River, NJ.

Prajapat M, Köhler J, Turchetta M, Krause A and Zeilinger MN (2025) Safe guaranteed exploration for non-linear systems. *IEEE Transactions on Automatic Control* .

Proctor JL, Brunton SL and Kutz JN (2016) Dynamic mode decomposition with control. *SIAM Journal on Applied Dynamical Systems* 15(1): 142–161.

Rasmussen DA, Ratmann O and Koelle K (2011) Inference for nonlinear epidemiological models using genealogies and time series. *PLoS computational biology* 7(8): e1002136.

12 Appendix A : Index to multimedia extensions

Table 2. Table of Multimedia Extension

Extension	Media Type	Description
1	Video	Comparative Analysis of SACK For Robotic Systems

Sah CK, Singh R and Keshavan J (2024) Real-time constrained tracking control of redundant manipulators using a koopman-zeroing neural network framework. *IEEE Robotics and Automation Letters* 9(2): 1732–1739.

Selim M, Bhat S and Johansson KH (????) Metakoopman: Bayesian meta-learning of koopman operators for modeling structured dynamics under distribution shifts. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Singh R, Sah CK and Keshavan J (2025a) Adaptive koopman embedding for robust control of nonlinear dynamical systems. *The International Journal of Robotics Research* 44(13): 2235–2261.

Singh R, Singh A, Kashyap CS and Keshavan J (2025b) Generalized momenta-based koopman formalism for robust control of euler-lagrangian systems. *arXiv preprint arXiv:2509.17010* .

Soloperto R, Köhler J and Allgöwer F (2020) Augmenting mpc schemes with active learning: Intuitive tuning and guaranteed performance. *IEEE Control Systems Letters* 4(3): 713–718.

Strogatz SH (2018) *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. CRC press.

Uchida D and Duraisamy K (2024) Model predictive control of nonlinear dynamics using online adaptive koopman operators. *arXiv preprint arXiv:2412.02972* .

Verschueren R, Frison G, Kouzoupis D, Frey J, van Duijkeren N, Zanelli A, Novoselnik B, Albin T, Quirynen R and Diehl M (2021) acados – a modular open-source framework for fast embedded optimal control. *Mathematical Programming Computation* .

Vovk V, Gammerman A and Shafer G (2005) *Algorithmic Learning in a Random World*. New York, NY: Springer. ISBN 978-0-387-00152-4. DOI:10.1007/b106715.

Wang Z, Lam J, Wei G, Fraser K and Liu X (2008) Filtering for nonlinear genetic regulatory networks with stochastic disturbances. *IEEE Transactions on Automatic Control* 53(10): 2448–2457.

Yang CM and Bhounsule PA (2025) Koopman operator based linear model predictive control for quadruped trotting. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 12359–12364.

Zinage V and Bakolas E (2021) Koopman operator based modeling for quadrotor control on se (3). *IEEE Control Systems Letters* 6: 752–757.

13 Appendix B - Offline Training and Implementation Details

13.1 Neural Network-Based Offline Learning

The offline training phase produces the three quantities that the online adaptive module requires: the lifting map $\psi(\cdot)$, the nominal Koopman matrices $(\boldsymbol{A}, \boldsymbol{B})$, and the reconstruction matrix $\boldsymbol{C}$. These are learned jointly from a nominal input-output dataset $\mathcal{D}_{\text{nom}} = \{(\boldsymbol{X}_i, \boldsymbol{Y}_i, \boldsymbol{U}_i)\}_{i=1}^{M}$, where $\boldsymbol{X}_i$ and $\boldsymbol{Y}_i$ are matrices of consecutive state observations and $\boldsymbol{U}_i$ the corresponding inputs, collected under nominal operating conditions. Once training is complete, $\psi(\cdot)$, and $\boldsymbol{C}$ are held fixed for the remainder of deployment. The quality of the offline-learned $\psi(\cdot)$ determines the validity of Assumption 2: if the nominal training data provides adequate coverage of the observable subspace, the perturbed dynamics will remain representable within the span of $\{\phi_1, \ldots, \phi_p\}$, and online adaptation over $(\boldsymbol{A}, \boldsymbol{B})$ alone will be sufficient to compensate for distributional shift.

We adopt the autoencoder-based Koopman architecture of Singh et al. (2025a), illustrated in Figure 1. In this architecture, the encoder network realizes the lifting map $\boldsymbol{z}_k = \psi(\boldsymbol{x}_k) \in \mathbb{R}^p$, a linear layer learns the Koopman operator matrices $\boldsymbol{A}$ and $\boldsymbol{B}$ governing the lifted dynamics $\boldsymbol{z}_{k+1} = \boldsymbol{A}\boldsymbol{z}_k + \boldsymbol{B}\boldsymbol{u}_k$, and a linear decoder realizes the reconstruction $\hat{\boldsymbol{x}}_k = \boldsymbol{C}\boldsymbol{z}_k$. The autoencoder structure is natural for this problem because it enforces the Koopman requirement that the lifted state be both dynamically consistent, evolving linearly under $(\boldsymbol{A}, \boldsymbol{B})$, and physically interpretable, decodable back to the original state space via $\boldsymbol{C}$. The network is trained end-to-end by minimizing a composite loss that encodes the three structural requirements of a valid Koopman representation:

$$\begin{aligned}\mathcal{L}_{\text{nom}} = \alpha_1\mathcal{L}_{\text{rec}} + \alpha_2\mathcal{L}_{\text{pred}} + \alpha_3\mathcal{L}_{\text{lift}} \\ + \gamma_1\|\boldsymbol{\Theta}\|_1 + \gamma_2\|\boldsymbol{\Theta}\|_2,\end{aligned} \tag{60}$$

where $\boldsymbol{\Theta}$ denotes all trainable network parameters and $\alpha_1, \alpha_2, \alpha_3, \gamma_1, \gamma_2 > 0$ are weighting hyperparameters. The three loss terms are: $\mathcal{L}_{\text{rec}} = \left\|\boldsymbol{x}_k - \boldsymbol{C}\boldsymbol{z}_k\right\|^2$, $\mathcal{L}_{\text{pred}} = \left\|\boldsymbol{x}_{k+1} - \boldsymbol{C}\hat{\boldsymbol{z}}_{k+1|k}\right\|^2$ and $\mathcal{L}_{\text{lift}} = \left\|\boldsymbol{z}_{k+1} - \hat{\boldsymbol{z}}_{k+1|k}\right\|^2$, where $\hat{\boldsymbol{z}}_{k+1|k} := \boldsymbol{A}\boldsymbol{z}_k + \boldsymbol{B}\boldsymbol{u}_k$ is the one-step linear prediction in the lifted space. $\mathcal{L}_{\text{rec}}$ enforces that $\boldsymbol{C}$ is a valid decoder, i.e., that $\boldsymbol{C}\psi(\boldsymbol{x}) \approx \boldsymbol{x}$ for all $\boldsymbol{x}$ in the training distribution. $\mathcal{L}_{\text{pred}}$ enforces that the linear lifted dynamics predict the next physical state accurately. $\mathcal{L}_{\text{lift}}$ enforces the Koopman invariance condition directly in the lifted space; it is the loss term most tightly coupled to the requirement that $\text{span}\{\phi_1, \ldots, \phi_p\}$ be approximately invariant under the system flow. The L_1 and L_2 regularization terms penalize all trainable parameters $\boldsymbol{\Theta}$ to reduce overfitting. For ease of learning, we include the base state, $\boldsymbol{x}$ as a part of the lifted state, so $\boldsymbol{z} = [\boldsymbol{x}^\top,\ \psi_{n_x+1},\ \ldots,\ \psi_p]^\top$. In this case the $\boldsymbol{C}$ matrix simplifies to $\boldsymbol{C} = [\boldsymbol{I}_{n_x \times n_x}^\top \boldsymbol{0}_{n_x \times (p-n_x)}^\top]^\top$ and the loss $\mathcal{L}_{\text{rec}}$ becomes redundant. The specific values of all hyperparameters for different robotic platforms are tabulated in Table 3.

Upon convergence, offline training yields the tuple $\{\psi(\cdot), \boldsymbol{A}, \boldsymbol{B}, \boldsymbol{C}\}$. The matrices $(\boldsymbol{A}, \boldsymbol{B})$ initialize the adaptive estimates: $\hat{\boldsymbol{A}}_0 = \boldsymbol{A}$ and $\hat{\boldsymbol{B}}_0 = \boldsymbol{B}$, so that the initial parameter estimation error is $\boldsymbol{E}_0 = \boldsymbol{W}_0^* - \hat{\boldsymbol{W}}_0 = [\Delta\boldsymbol{A}_0,\ \Delta\boldsymbol{B}_0]$. The magnitude $\|\boldsymbol{E}_0\|_F$ quantifies the initial mismatch between the nominal and true Koopman operators at deployment time and appears explicitly in the finite-time convergence bound of Theorem 1. A higher-quality offline model, one trained on data that better covers the deployment

distribution, yields a smaller $\|\boldsymbol{E}_0\|_F$ and therefore faster convergence of the online adaptation law.

13.2 Implementation Details

In all cases, the Koopman matrices $\boldsymbol{A}$, $\boldsymbol{B}$, $\boldsymbol{C}$ are identified from the collected data using a neural network architecture. The hyperparameters for the employed architecture are shown in Table 3. In Table 3, the architecture of the 3R manipulator is written as $[6, 30, 30, 17]$, which basically means that the encoder has an input layer of dimension 6 corresponding to each state, there are two hidden layers, each of width 30, and the output dimension is 17, which basically is the dimension of the lifted states. The same nomenclature holds for the other systems.

Table 3. Hyperparameters of the Koopman network and CBF implemetatoion.

	3R manipulator	Planar Quadrotor	Franka Research 3	TurtleBot 3
Nominal Koopman Autoencoder (linear)				
Architecture	[6, 30, 30, 17]	[6, 20, 20, 17]	[14, 30, 30, 73]	[3, 30, 30, 14]
# lifted state	17	17	73	14
α_1, α_2, α_3	1, 0.3, 1	1, 0.3, 1	1, 0.5, 1	1, 0.3, 1
γ_1, γ_2	0, 0	0, 0	0, 0	0, 0
Batch Size	256	256	256	256
CBF Implementation				
n_{conf}	40	40	40	40
χ	0.025	0.025	0.025	0.025
α_{cbf}	0.1	0.1	0.1	0.1
α_{EMA}	0.2	0.2	0.2	0.2

13.3 Data Generation

Nominal Koopman models are identified offline from simulation data prior to deployment. For the 3R serial manipulator, planar quadrotor, and Franka Research 3, training trajectories are generated using minimum-snap trajectory optimization through randomly sampled waypoints. Segment timing is optimized under a maximum velocity constraint, and 10^{th}-order polynomial coefficients are solved via unconstrained endpoint optimization, producing C^4-continuous reference trajectories that excite a broad range of configurations and velocities. The system is simulated under these references and state-input pairs are logged at the control frequency.

For the TurtleBot3 Burger, velocity command trajectories are collected in Gazebo under a combination of randomized inputs and sinusoidal velocity profiles spanning the full admissible range $\boldsymbol{v} \in [-0.176, 0.176]$ m/s, $\omega \in [-2.272, 2.272]$ rad/s. The sinusoidal profiles ensure smooth, persistently exciting trajectories that cover the nonlinear kinematic regime, while the random inputs provide broad coverage of the input space.

14 Appendix C- Proofs for Section 4

14.1 Proof of Lemma 2

The rate in Lemma 2 is

$$\rho_T := \Big(1 - \frac{(2-q)\,\eta\alpha}{2 + q^2 T(T-1)}\Big)^{1/2} < 1, \tag{61}$$

with $q := \eta\bar{\lambda} \in (0,2)$, and, writing $k = mT + r$ with $0 \le r < T$, the resulting bound on the estimation error is

$$\|\boldsymbol{E}_k\|_F \le \rho_T^m \|\boldsymbol{E}_0\|_F + \frac{1-\rho_T^m}{1-\rho_T}\, T\nu_{\text{eff}} + r\,\nu_{\text{eff}}. \tag{62}$$

Proof of Lemma 2. Unrolling (21) gives $\boldsymbol{E}_{k+T} = \boldsymbol{E}_k\boldsymbol{P} + \sum_{j=k}^{k+T-1} \tilde{\boldsymbol{\Delta}}_j(\boldsymbol{I} - \eta\boldsymbol{G}_{j+1})\cdots(\boldsymbol{I} - \eta\boldsymbol{G}_{k+T-1})$ with $\boldsymbol{P} := (\boldsymbol{I} - \eta\boldsymbol{G}_k)\cdots(\boldsymbol{I} - \eta\boldsymbol{G}_{k+T-1})$. Every factor has spectral norm at most one by Lemma 1(i), so the sum is bounded by $T\nu_{\text{eff}}$, and it remains to show $\|\boldsymbol{P}\|_2 \le \rho_T$. Let $q := \eta\bar{\lambda} \in (0,2)$ and $\boldsymbol{M}_i := \eta\boldsymbol{G}_{k+i}$, so $\boldsymbol{0} \preceq \boldsymbol{M}_i \preceq q\boldsymbol{I}$, and for a unit vector $\boldsymbol{x}_0$ let $\boldsymbol{x}_{i+1} := (\boldsymbol{I} - \boldsymbol{M}_i)\boldsymbol{x}_i$, so that $\boldsymbol{x}_T = \boldsymbol{P}^\top\boldsymbol{x}_0$. Since $2\boldsymbol{M}_i - \boldsymbol{M}_i^2 \succeq (2-q)\boldsymbol{M}_i$, $\|\boldsymbol{x}_{i+1}\|_2^2 \le \|\boldsymbol{x}_i\|_2^2 - (2-q)\boldsymbol{x}_i^\top\boldsymbol{M}_i\boldsymbol{x}_i$, hence $\|\boldsymbol{x}_T\|_2^2 \le 1 - (2-q)\Sigma$ with $\Sigma := \sum_{i=0}^{T-1}\boldsymbol{x}_i^\top\boldsymbol{M}_i\boldsymbol{x}_i$. Moreover $\boldsymbol{x}_0 - \boldsymbol{x}_i = \sum_{l<i}\boldsymbol{M}_l\boldsymbol{x}_l$ and $\|\boldsymbol{M}_l\boldsymbol{x}_l\|_2 \le \sqrt{q}\,(\boldsymbol{x}_l^\top\boldsymbol{M}_l\boldsymbol{x}_l)^{1/2}$, so the Cauchy–Schwarz inequality gives $\|\boldsymbol{x}_0 - \boldsymbol{x}_i\|_2^2 \le q\,i\,\Sigma$. Using $(\boldsymbol{a}+\boldsymbol{b})^\top\boldsymbol{M}(\boldsymbol{a}+\boldsymbol{b}) \le 2\boldsymbol{a}^\top\boldsymbol{M}\boldsymbol{a} + 2\boldsymbol{b}^\top\boldsymbol{M}\boldsymbol{b}$ with $\boldsymbol{a} = \boldsymbol{x}_i$, $\boldsymbol{b} = \boldsymbol{x}_0 - \boldsymbol{x}_i$, and (24),

$$\begin{aligned}\eta\alpha \le \sum_{i=0}^{T-1}\boldsymbol{x}_0^\top\boldsymbol{M}_i\boldsymbol{x}_0 &\le 2\Sigma + 2q\sum_{i=0}^{T-1}\|\boldsymbol{x}_0 - \boldsymbol{x}_i\|_2^2 \\ &\le \big(2 + q^2T(T-1)\big)\Sigma.\end{aligned}$$

Hence $\|\boldsymbol{P}^\top\boldsymbol{x}_0\|_2^2 \le \rho_T^2$ for every unit $\boldsymbol{x}_0$. The radicand of ρ_T is positive because $\eta\alpha \le qT$ and $2 + q^2T(T-1) - (2-q)qT = (qT-1)^2 + 1 > 0$. Applying the block bound m times and Lemma 1(i) for the remaining r steps yields (62).

14.2 Realized excitation

Corollary 1. Realized excitation. Let $\boldsymbol{U} = \boldsymbol{U}_k^* \in \mathcal{F}$ be the input sequence computed at time k, whose first element is applied, let $\boldsymbol{v}_{k+i}$, $i = 0, \dots, N_p - 1$, denote the regressors realized in closed loop, and suppose $\|\boldsymbol{v}_{k+i} - \hat{\boldsymbol{v}}_{k+i|k}\|_2 \le \vartheta_i$ and $\|\boldsymbol{v}_{k+i}\|_2 \le v_{\max}$. Then

$$\lambda_{\min}\Big(\sum_{j=k}^{k+N_p-1}\boldsymbol{G}_j\Big) \;\ge\; \lambda_{\min}\big(\boldsymbol{S}_k(\boldsymbol{U})\big) - 2v_{\max}\sum_{i=1}^{N_p-1}\vartheta_i. \tag{63}$$

Proof Corollary 1. Since $\hat{\boldsymbol{z}}_{k|k} = \boldsymbol{z}_k$ by (29c) and $\boldsymbol{u}_k = \boldsymbol{u}_{k|k}$, $\vartheta_0 = 0$. Each stored regressor $\boldsymbol{v}_{k-m}$, $1 \le m \le w-1$, enters $\boldsymbol{G}_k$ with weight γ^m, and each realized $\boldsymbol{v}_{k+i}$ enters $\boldsymbol{G}_{k+i}$ with weight one; all other terms are positive semidefinite. Hence $\sum_{j=k}^{k+N_p-1}\boldsymbol{G}_j \succeq \sum_{m=1}^{w-1}\gamma^m\boldsymbol{v}_{k-m}\boldsymbol{v}_{k-m}^\top + \sum_{i=0}^{N_p-1}\boldsymbol{v}_{k+i}\boldsymbol{v}_{k+i}^\top$. The right-hand side differs from $\boldsymbol{S}_k(\boldsymbol{U})$ by $\sum_{i\ge1}(\boldsymbol{v}_{k+i}\boldsymbol{v}_{k+i}^\top - \hat{\boldsymbol{v}}_{k+i|k}\hat{\boldsymbol{v}}_{k+i|k}^\top)$, and $\|\boldsymbol{a}\boldsymbol{a}^\top - \boldsymbol{b}\boldsymbol{b}^\top\|_2 = \|\boldsymbol{a}(\boldsymbol{a}-\boldsymbol{b})^\top + (\boldsymbol{a}-\boldsymbol{b})\boldsymbol{b}^\top\|_2 \le 2v_{\max}\|\boldsymbol{a}-\boldsymbol{b}\|_2$ for $\|\boldsymbol{a}\|_2, \|\boldsymbol{b}\|_2 \le v_{\max}$. Weyl's inequality completes the proof.

15 Appendix D- Proofs of the Safety Guarantees

15.1 Shifted-candidate feasibility

Lemma 3. Shifted-candidate feasibility. Let the MPC at time k be feasible with optimal input sequence $\boldsymbol{U}_k^* = \{\hat{\boldsymbol{u}}_{k|k}^*, \dots, \hat{\boldsymbol{u}}_{k+N_p-1|k}^*\}$ and nominal prediction $\{\hat{\boldsymbol{z}}_{k+i|k}^*\}_{i=0}^{N_p}$, and let the first input $\boldsymbol{u}_k = \hat{\boldsymbol{u}}_{k|k}^*$ be applied. The measured re-initialization (29c) then gives

$$\hat{\boldsymbol{z}}_{k+1|k+1} = \boldsymbol{z}_{k+1} = \hat{\boldsymbol{z}}_{k+1|k}^* + \boldsymbol{d}_k, \ \|\boldsymbol{d}_k\|_2 \le \bar{\delta}^{\mathrm{ana}}. \quad (64)$$

Define the candidate sequence for time $k+1$ by

$$\begin{aligned} \hat{\boldsymbol{u}}_{k+1+m|k+1} &:= \hat{\boldsymbol{u}}_{k+1+m|k}^* + \boldsymbol{K}_{\mathrm{ctrl}}\boldsymbol{\zeta}_m, \\ \hat{\boldsymbol{z}}_{k+1+m|k+1} &:= \hat{\boldsymbol{z}}_{k+1+m|k}^* + \boldsymbol{\zeta}_m, \end{aligned} \quad (65)$$

for $m = 0, \dots, N_p - 2$, where $\boldsymbol{\zeta}_m$ is generated by

$$\begin{aligned} \boldsymbol{\zeta}_0 &:= \boldsymbol{d}_k, \qquad \boldsymbol{\zeta}_{m+1} = \boldsymbol{A}_{\mathrm{cl},k+1}\boldsymbol{\zeta}_m + \boldsymbol{\eta}_m, \\ \boldsymbol{\eta}_m &:= (\hat{\boldsymbol{A}}_{k+1} - \hat{\boldsymbol{A}}_k)\,\hat{\boldsymbol{z}}_{k+1+m|k}^* \\ &\quad + (\hat{\boldsymbol{B}}_{k+1} - \hat{\boldsymbol{B}}_k)\,\hat{\boldsymbol{u}}_{k+1+m|k}^*. \end{aligned} \quad (66)$$

Then:

(i) the candidate (65) satisfies the prediction dynamics (29b) under the updated model $(\hat{\boldsymbol{A}}_{k+1}, \hat{\boldsymbol{B}}_{k+1})$ and the initialization (64);

(ii) $\|\boldsymbol{\eta}_m\|_2 \le \bar{\mu}\, v_{\max}$ and $\|\boldsymbol{\zeta}_m\|_2 \le b_m$, $m = 0, \dots, N_p - 1$, with b_m from (45);

(iii) the tightened sets (46) satisfy: if $\hat{\boldsymbol{z}}_{k+1+m|k}^* \in \mathcal{Z}_{m+1}^t$ and $\hat{\boldsymbol{u}}_{k+1+m|k}^* \in \mathcal{U}_{m+1}^t$, then $\hat{\boldsymbol{z}}_{k+1+m|k+1} \in \mathcal{Z}_m^t$ and $\hat{\boldsymbol{u}}_{k+1+m|k+1} \in \mathcal{U}_m^t$, for $m = 0, \dots, N_p - 2$.

Proof. (i) Under $(\hat{\boldsymbol{A}}_{k+1}, \hat{\boldsymbol{B}}_{k+1})$, the candidate evolves as $\hat{\boldsymbol{A}}_{k+1}(\hat{\boldsymbol{z}}_{k+1+m|k}^* + \boldsymbol{\zeta}_m) + \hat{\boldsymbol{B}}_{k+1}(\hat{\boldsymbol{u}}_{k+1+m|k}^* + \boldsymbol{K}_{\mathrm{ctrl}}\boldsymbol{\zeta}_m) = \hat{\boldsymbol{A}}_{k+1}\hat{\boldsymbol{z}}_{k+1+m|k}^* + \hat{\boldsymbol{B}}_{k+1}\hat{\boldsymbol{u}}_{k+1+m|k}^* + \boldsymbol{A}_{\mathrm{cl},k+1}\boldsymbol{\zeta}_m$. Subtracting the time-$k$ prediction $\hat{\boldsymbol{z}}_{k+2+m|k}^* = \hat{\boldsymbol{A}}_k\hat{\boldsymbol{z}}_{k+1+m|k}^* + \hat{\boldsymbol{B}}_k\hat{\boldsymbol{u}}_{k+1+m|k}^*$ shows the candidate state at stage $m+1$ equals $\hat{\boldsymbol{z}}_{k+2+m|k}^* + \boldsymbol{\zeta}_{m+1}$ with $\boldsymbol{\zeta}_{m+1}$ as in (66), i.e., the candidate is dynamically consistent; the base case is (64) with $\boldsymbol{\zeta}_0 = \boldsymbol{d}_k$.

(ii) By (38), $\|[\hat{\boldsymbol{A}}_{k+1} - \hat{\boldsymbol{A}}_k, \hat{\boldsymbol{B}}_{k+1} - \hat{\boldsymbol{B}}_k]\|_F \le \mu_k \le \bar{\mu}$, and the stacked nominal regressor lies in $\mathcal{Z} \times \mathcal{U}$, so $\|\boldsymbol{\eta}_m\|_2 \le \bar{\mu}\, v_{\max}$. The model is fixed over the horizon at time $k+1$, so $\boldsymbol{A}_{\mathrm{cl},k+1}$ is constant in the recursion, and (42) applies with the constant sequence: unrolling (66),

$$\boldsymbol{\zeta}_m = \boldsymbol{A}_{\mathrm{cl},k+1}^m \boldsymbol{d}_k + \sum_{\ell=0}^{m-1} \boldsymbol{A}_{\mathrm{cl},k+1}^{m-1-\ell}\boldsymbol{\eta}_\ell,$$

$$\|\boldsymbol{\zeta}_m\|_2 \le c\tilde{\alpha}^{\,m}\bar{\delta}^{\mathrm{ana}} + \sum_{\ell=0}^{m-1} c\tilde{\alpha}^{\,m-1-\ell}\bar{\mu}v_{\max} = b_m.$$

(iii) By construction $\varpi_{m+1} = \varpi_m + b_m$, hence $\mathcal{B}(b_m) \oplus \mathcal{B}(\varpi_m) \subseteq \mathcal{B}(\varpi_{m+1})$, and by monotonicity of the Pontryagin difference, $\hat{\boldsymbol{z}}_{k+1+m|k+1} = \hat{\boldsymbol{z}}_{k+1+m|k}^* + \boldsymbol{\zeta}_m \in \mathcal{Z}_{m+1}^t \oplus \mathcal{B}(b_m) \subseteq \mathcal{Z}_m^t$. The input claim is identical with $\|\boldsymbol{K}_{\mathrm{ctrl}}\boldsymbol{\zeta}_m\|_2 \le \|\boldsymbol{K}_{\mathrm{ctrl}}\|_2 b_m$.

Remark 5. The candidate correction absorbs the two perturbations that re-solving introduces: the realized disturbance $\boldsymbol{d}_k$, which enters through the measured re-initialization (64) and decays along the horizon at rate $\tilde{\alpha}$, and the model update $(\hat{\boldsymbol{W}}_{k+1} - \hat{\boldsymbol{W}}_k)$, which accumulates through $\boldsymbol{\eta}_m$ toward the steady offset $c\bar{\mu}v_{\max}/(1-\tilde{\alpha})$. The feedback gain $\boldsymbol{K}_{\mathrm{ctrl}}$ appears only in the *candidate* construction; the implemented input remains $\boldsymbol{u}_k = \hat{\boldsymbol{u}}_{k|k}^*$, consistent with the re-initialization (29c).

15.2 Proof of Theorem 3

Proof of Theorem 3.

(i) Suppose (29) is feasible at time k with optimizer $\boldsymbol{U}_k^*$ and nominal trajectory satisfying $\hat{\boldsymbol{z}}_{k+i|k}^* \in \mathcal{Z}_i^t$, $\hat{\boldsymbol{u}}_{k+i|k}^* \in \mathcal{U}_i^t$ for $i \le N_p - 1$ and $\hat{\boldsymbol{z}}_{k+N_p|k}^* \in \mathcal{Z}_f$; this holds at $k=0$ by Assumption 9. Construct the candidate of Lemma 3. By Lemma 3(i) it satisfies the prediction dynamics and re-initialization, and by Lemma 3(iii) its stages $m = 0, \dots, N_p - 2$ satisfy the tightened state and input constraints. For the terminal stage, the candidate state at stage $N_p - 1$ is $\tilde{\boldsymbol{z}} := \hat{\boldsymbol{z}}_{k+N_p|k}^* + \boldsymbol{\zeta}_{N_p-1} \in \mathcal{Z}_f \oplus \mathcal{B}(\bar{\delta}^+)$, since $b_m \le \bar{\delta}^+$ for all m by (45) and (43); the additional containments in the theorem statement place $\tilde{\boldsymbol{z}} \in \mathcal{Z}_{N_p-1}^t$ and the terminal candidate input $\kappa_f(\tilde{\boldsymbol{z}}) = \boldsymbol{K}_{\mathrm{ctrl}}\tilde{\boldsymbol{z}} \in \mathcal{U}_{N_p-1}^t$, while Assumption 7(ii) with $\hat{\boldsymbol{W}}_{k+1} \in \mathcal{W}$ gives $(\hat{\boldsymbol{A}}_{k+1} + \hat{\boldsymbol{B}}_{k+1}\boldsymbol{K}_{\mathrm{ctrl}})\tilde{\boldsymbol{z}} \in \mathcal{Z}_f$.

For the CBF rows, write $\tilde{\boldsymbol{x}}_m := \boldsymbol{C}(\hat{\boldsymbol{z}}_{k+1+m|k}^* + \boldsymbol{\zeta}_m)$. Assumption 8 gives $|h(\tilde{\boldsymbol{x}}_m) - h(\boldsymbol{C}\hat{\boldsymbol{z}}_{k+1+m|k}^*)| \le L_h\|\boldsymbol{C}\|_2 b_m$, so the time-$k$ row at stage $m+1$ yields $h(\tilde{\boldsymbol{x}}_{m+1}) \ge (1-\alpha_{\mathrm{cbf}})h(\tilde{\boldsymbol{x}}_m) + L_h\sigma_{m+1} - L_h\|\boldsymbol{C}\|_2\big(b_{m+1} + (1-\alpha_{\mathrm{cbf}})b_m\big) = (1-\alpha_{\mathrm{cbf}})h(\tilde{\boldsymbol{x}}_m) + L_h\sigma_m$, for $m = 0, \dots, N_p - 2$, using $\sigma_{m+1} - \sigma_m = \|\boldsymbol{C}\|_2\big(b_{m+1} + (1-\alpha_{\mathrm{cbf}})b_m\big)$ from (47) and $\varpi_{j+1} - \varpi_j = b_j$. The last row, $m = N_p - 1$, connects $\tilde{\boldsymbol{z}}$ to the new terminal state $(\hat{\boldsymbol{A}}_{k+1} + \hat{\boldsymbol{B}}_{k+1}\boldsymbol{K}_{\mathrm{ctrl}})\tilde{\boldsymbol{z}}$ and has no counterpart in the time-k plan; it holds by (48), because $\hat{\boldsymbol{W}}_{k+1} \in \mathcal{W}$, $\tilde{\boldsymbol{z}} \in \mathcal{Z}_f \oplus \mathcal{B}(\bar{\delta}^+)$, and the stage margin σ_{N_p-1} at time $k+1$ is at most $\bar{\sigma}$, since $\delta_{x,k+1}^{\mathrm{ana}} \le \|\boldsymbol{C}\|_2\bar{\delta}^{\mathrm{ana}}$. The margins at time $k+1$ use $\delta_{x,k+1}^{\mathrm{ana}} \le \delta_{x,k}^{\mathrm{ana}}$, since $\bar{E}_k$ in (35) is non-increasing, so the candidate also satisfies the time-$(k+1)$ rows. Hence the candidate satisfies every constraint of the time-$(k+1)$ problem, completing the induction.

(ii) The stage-0 row at time k gives $h(\hat{\boldsymbol{x}}_{k+1|k}) \ge (1-\alpha_{\mathrm{cbf}})h(\boldsymbol{x}_k) + L_h\sigma_0 \ge (1-\alpha_{\mathrm{cbf}})h(\boldsymbol{x}_k) + L_h\delta_{x,k}^{\mathrm{ana}}$. Since $\boldsymbol{x}_{k+1} = \hat{\boldsymbol{x}}_{k+1|k} + \boldsymbol{C}\boldsymbol{d}_k$ with $\|\boldsymbol{C}\boldsymbol{d}_k\|_2 \le \delta_{x,k}^{\mathrm{ana}}$, Assumption 8 yields $h(\boldsymbol{x}_{k+1}) \ge h(\hat{\boldsymbol{x}}_{k+1|k}) - L_h\delta_{x,k}^{\mathrm{ana}} \ge (1-\alpha_{\mathrm{cbf}})h(\boldsymbol{x}_k)$, where the first inequality is Assumption 8 and the second is the MPC constraint. Since $h(\boldsymbol{x}_0) \ge 0$, induction gives $\boldsymbol{x}_k \in \mathcal{S}$ for all k.

(iii) The first claim is (64) with $\|\boldsymbol{d}_k\|_2 \le \delta_k^{\mathrm{ana}}$ and the limit superior of $\bar{E}_k$ in (35).

15.3 Proof of Theorem 4

Proof of Theorem 4. **(i)** Under Assumption 10, the rank of r_k among $\{r_{k-n_{conf}}, \dots, r_{k-1}, r_k\}$ is uniformly distributed over $\{1, \dots, n_{conf}+1\}$ (with ties broken uniformly at random). Since $\hat{Q}_{1-\chi}$ is the $\lceil (n_{conf}+1)(1-\chi)\rceil$-th smallest of the n_{conf} calibration scores, r_k exceeds it only if its rank is larger than this index, so, with $\mathcal{R}_k := \{r_j : j \in$

$\mathcal{W}_k\}$ Vovk et al. (2005),

$$\begin{aligned}\mathbb{P}\big(r_k \le \hat{Q}_{1-\chi}(\mathcal{R}_k)\big) &\ge \frac{\lceil (n_{conf}+1)(1-\chi)\rceil}{n_{conf}+1} \\ &\ge 1-\chi.\end{aligned} \tag{67}$$

Define the event $\mathcal{E}_k := \big\{r_k \le \hat{Q}_{1-\chi}(\{r_j : j \in \mathcal{W}_k\})\big\}$. On $\mathcal{E}_k$, the triangle inequality and the definition (51) of the score give

$$\begin{aligned}\|\boldsymbol{C}\boldsymbol{d}_k\|_2 = \|\boldsymbol{s}_k\|_2 &\le \big\|\boldsymbol{s}_k - \hat{\boldsymbol{s}}_{k-1}\big\|_2 + \big\|\hat{\boldsymbol{s}}_{k-1}\big\|_2 \\ &= r_k + \big\|\hat{\boldsymbol{s}}_{k-1}\big\|_2 \;\le\; \delta_k^{\text{conf}},\end{aligned} \tag{68}$$

where the last inequality uses (52). Note that (68) holds without any assumption on the magnitude of the EMA, since $\|\hat{\boldsymbol{s}}_{k-1}\|_2$ enters the margin as a measured quantity.

Applying Assumption 8 with $\boldsymbol{e} = \boldsymbol{C}\boldsymbol{d}_k$ and substituting (68),

$$\begin{aligned}h(\boldsymbol{x}_{k+1}) &\ge h\big(\hat{\boldsymbol{x}}_{k+1|k}\big) - L_h\|\boldsymbol{C}\boldsymbol{d}_k\|_2 \\ &\ge h\big(\hat{\boldsymbol{x}}_{k+1|k}\big) - L_h\,\delta_k^{\text{conf}}.\end{aligned} \tag{69}$$

Substituting (55) into (69) on $\mathcal{E}_k$,

$$\begin{aligned}h(\boldsymbol{x}_{k+1}) &\ge (1-\alpha_{cbf})\,h(\boldsymbol{x}_k) + L_h\,\delta_k^{\text{conf}} - L_h\,\delta_k^{\text{conf}} \\ &\ge (1-\alpha_{cbf})\,h(\boldsymbol{x}_k).\end{aligned} \tag{70}$$

Since this holds on $\mathcal{E}_k$ and $P(\mathcal{E}_k) \ge 1-\chi$ by (67), (56) follows.

(ii) Define the failure event at each step as $\mathcal{F}_k := \big\{h(\boldsymbol{x}_{k+1}) < (1-\alpha_{cbf})\,h(\boldsymbol{x}_k)\big\}$. From (i), $P(\mathcal{F}_k) \le \chi$ for each $k \ge k_{warm}$. By the union bound,

$$P\left(\bigcup_{k=k_{warm}}^{k_{warm}+H-1} \mathcal{F}_k\right) \le \sum_{k=k_{warm}}^{k_{warm}+H-1} P(\mathcal{F}_k) \le H\chi. \tag{71}$$

On the complement event, the CBF evolution condition holds at every step, so $h(\boldsymbol{x}_k) \ge (1-\alpha_{cbf})^{k-k_{warm}} h(\boldsymbol{x}_{k_{warm}}) \ge 0$ for all $k \in \{k_{warm}+1, \ldots, k_{warm}+H\}$, establishing (57).

Theorem 4 applies to δ_k^{warm} unchanged, with (68) replaced by the direct bound $\|\boldsymbol{C}\boldsymbol{d}_k\|_2 = \|\boldsymbol{s}_k\|_2 \le \delta_k^{warm}$ on the corresponding coverage event, since the warm-up scores are the raw disturbance norms and require no bias correction.