MINUTES OF DATA, MILLISECONDS OF CONTROL: SHARED PHYSICAL STRUCTURE FOR QUADROTOR WORLD MODELING AND CONVEX MPC

Anonymous authors
Paper under double-blind review

ABSTRACT

Quadrotor world models must learn from limited noisy observations while supporting optimization at high control rates. We introduce ABS-PWM, a probabilistic world model based on an analytically constructed bilinear lifting of quadrotor dynamics. Its analytical lifting and reconstruction maps maintain physically meaningful predictions, while a structured stochastic transition captures aerodynamic effects, unmodeled dynamics, and observation uncertainty. We train ABS-PWM directly from noisy trajectories using generalized expectation conditional maximization, combining latent-state smoothing with closed-form covariance updates. The same bilinear structure enables BUC-MPC, a convex uncertainty-aware model predictive control (MPC) scheme. Experiments in high-fidelity simulation show that ABS-PWM achieves competitive short-horizon prediction accuracy from limited flight data while providing faster rollouts than the evaluated state-of-the-art baselines. In closed loop, BUC-MPC achieves accurate tracking, high completion rates under shifts in vehicle and environmental dynamics, and computation compatible with high-frequency control. Additional experiments with a separately trained model on a second quadrotor platform demonstrate the framework's applicability across vehicle configurations.

1 INTRODUCTION

Quadrotors support diverse applications, including logistics, communications, inspection, search and rescue, surveillance, and agriculture [1]–[4]. Quadrotors' nonlinear, underactuated dynamics and high-frequency control requirements have drawn considerable research attention [5]–[6]. An effective quadrotor controller should accurately track trajectories while anticipating future references, enforce system constraints, remain robust to disturbances and unmodeled dynamics, generalize well across tasks or trajectories, and operate in real-time despite the fast dynamics of the vehicle [7]–[8]. Predictive control methods such as MPC offer many of these desirable features for quadrotor control [9], although computationally expensive nonconvex optimization can challenge real-time implementation. Their effectiveness also depends on the accuracy of the underlying dynamics model.

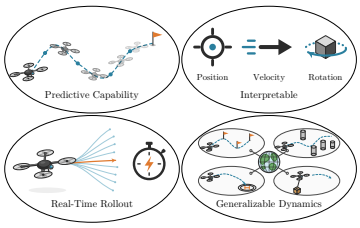


Figure 1: Desirable features of a quadrotor WM

Classical physics-based models provide interpretable descriptions of quadrotor dynamics and remain effective near nominal operating conditions. However, tractable physics-based models often simplify aerodynamic effects that become significant during agile flight, degrading prediction and tracking accuracy [10]–[11]. Their reliability further depends on precisely identified parameters (e.g., mass, inertia, drag) that fluctuate with payload and hardware, alongside unpredictable external disturbances like wind [12]–[13]. Higher-fidelity models can reduce these discrepancies but require substantial modeling and experimental effort and may be costly for real-time prediction [11]–[12]. Balancing predictive fidelity and computational efficiency across vehicles and flight conditions therefore motivates a control-oriented world model (WM) that combines physical knowledge with dynamics learned from data [8].

A WM can serve as an internal simulator that captures the quadrotor and environment dynamics, enabling forward prediction and counterfactual rollouts for informed decision-making. Therefore,

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

a control-oriented WM for quadrotors should ideally provide (see Fig. 1) accurate and physically meaningful predictions, represent task-relevant physical quantities, support real-time rollouts, and learn a generalizable dynamics representation reusable across tasks and control architectures [8], [14]. Two requirements are particularly important for deployment: data efficiency and closed-loop computational efficiency. Learning from flight observations can help account for disturbances and changing operating conditions [15]–[16], but collecting representative data across vehicles, payloads, and environments is costly and can pose safety risks. This motivates data-efficient world-model learning [17]–[18] and the use of physical structure to further reduce data requirements. Meanwhile, fast forward prediction alone does not ensure efficient predictive control: the model must also support tractable optimization over future control inputs. These considerations motivate our central question: *How can a shared physical structure enable both probabilistic WM learning from limited noisy observations and computationally efficient predictive control?*

The key contributions of this work are: **(i) Analytical bilinear formulation:** We derive a novel lifted bilinear representation of nominal quadrotor dynamics with analytical encoding and reconstruction and an explicit truncation residual bound, then incorporate structured drag and damping without additional lifted states. **(ii) Probabilistic identification:** We introduce ABS-PWM (**A**nalytical **B**ilinear **S**tructured-**P**robabilistic **W**orld **M**odel) and a generalized expectation conditional maximization (GECM) procedure to jointly identify physical parameters and noise statistics from noisy observations, with closed-form covariance updates and a conditional surrogate-improvement guarantee. **(iii) Efficient uncertainty-aware control:** We develop BUC-MPC (**B**ilinear **U**ncertainty-aware **C**onvex **M**odel **P**redictive **C**ontrol), combining a convex tracking approximation based on fixed guide trajectories with covariance-weighted disturbance compensation. **(iv) Limited-data validation:** Using 90 s of noisy flight data for the Skydio X2 from a high-fidelity MuJoCo Menagerie environment, we evaluate prediction and closed-loop tracking under in-distribution (ID) and out-of-distribution (OOD) conditions. The learned model supports 100-step rollouts in approximately 1.3 ms, while BUC-MPC meets the computation budget for 100 Hz control. We further evaluate the framework on a Crazyflie 2 platform, demonstrating its applicability to a distinct vehicle platform.

2 Related Work

World models for quadrotors. WMs for quadrotors have been investigated for model-based RL (MBRL) [19]–[23]. MAD [24] additionally investigates frozen world-model encoders as representations for downstream RL. For predictive control, SkyJEPA embeds learned latent dynamics within MPPI [8], whereas AeroAct generates trajectory references tracked by a separate MPC [25]. Other approaches use predictive objectives for offline navigation-policy learning [26] or motion-conditioned aerial video generation [27]. Despite these advances, existing quadrotor world models often rely on extensive simulated interactions or large offline datasets, motivating data-efficient dynamics learning with fast rollouts for real-time closed-loop control.

Neural ODE approaches. KNODE [28] augments first-principles dynamics with learned neural residuals for predictive control . Online extensions adapt these residual models during deployment to accommodate changing disturbances [15], while feedback neural ODEs [29] correct learned dynamics to improve generalization under uncertainty . These approaches improve model accuracy and adaptability, but retain nonlinear prediction models within generally nonconvex MPC problems, motivating structured alternatives that reduce rollout and online optimization costs.

Koopman operator-based approaches. Koopman-based methods provide simplified representations of nonlinear dynamical systems using observable functions of the state [30]–[33]. For quadrotors, KQ-LMPC exploits an analytical Koopman representation of nominal dynamics for efficient closed-loop control [34]. KoopNet jointly learns a function dictionary and bilinear dynamics for quadrotor MPC [35], while Jacobian-regularized dynamic-mode decomposition incorporates derivatives of an approximate prior model to improve sample efficiency [36]. These control-oriented formulations employ deterministic prediction models. Koopman-based error-state Kalman filtering accounts for noise in state estimation through IMU-driven propagation [37], but does not directly provide a dynamics model conditioned on control commands for predictive control. This motivates probabilistic learning of control-conditioned dynamics that accounts for environmental disturbances and noisy measurements while enabling efficient predictive rollouts.

Physics-based and learned dynamics models. Physics-based and learned models improve quadrotor prediction and control by capturing aerodynamic effects and other unmodeled dynamics. Faessler et al. incorporate linear rotor drag into a differentially flat model [38], while NeuroBEM combines blade-element momentum theory with learned aerodynamic residuals [11]. Gaussian-process residual models [39] and Real-Time Neural-MPC [40] integrate learned dynamics into predictive control. Physics-inspired temporal networks combine state history with physics constraints for dynamics learning and MPC [41], while sequential architectures and multi-step objectives mitigate compounding prediction errors [42]. Hyper Prediction Models predict time-varying model parameters over the control horizon to capture unmodeled dynamics [43]. Other approaches combine offline and online dynamics learning with MPC informed by aleatoric uncertainty [16], or use conditional diffusion models for dynamics prediction and adaptive control [44]. Nevertheless, balancing model expressiveness and uncertainty representation with the computational demands of repeated prediction and online control remains a central challenge.

3 PRELIMINARIES

Quadrotor nominal dynamics. Consider the inertial (world) reference frame $\mathcal{I}_{\text{ref}}$ and the body-fixed reference frame $\mathfrak{B}_{\text{ref}}$ of a quadrotor, depicted in Fig. 2 and explained in Definition 1. Let $\mathbf{s}, \mathbf{v} \in \mathbb{R}^3$ denote the position and linear velocity of the COM of the quadrotor expressed in $\mathcal{I}_{\text{ref}}$, respectively. Let $\mathbf{R} \in \mathrm{SO}(3)$ denote the rotation matrix from $\mathfrak{B}_{\text{ref}}$ to $\mathcal{I}_{\text{ref}}$. Furthermore, let $m > 0$ be the quadrotor mass, $\mathbf{J} \in \mathbb{R}^{3\times3}$ the inertia matrix in $\mathfrak{B}_{\text{ref}}$, $\boldsymbol{\omega} \in \mathbb{R}^3$ the angular velocity in $\mathfrak{B}_{\text{ref}}$, f the total thrust along $\vec{b}_3$, and $\boldsymbol{\tau} := [\tau_1,\ \tau_2,\ \tau_3]^\top \in \mathbb{R}^3$ the total moment expressed in $\mathfrak{B}_{\text{ref}}$. Assume that $\mathbf{J} = \mathbf{J}^\top \succ 0$, and define $\tilde{\mathbf{J}} := \mathbf{J}^{-1}$. The gravitational acceleration has magnitude g and we define $\vec{\mathbf{g}} = g\vec{e}_3$. For $i \in [1:3]$, let $\mathbf{S}_i := \vec{e}_i^{\wedge}$. Define $\boldsymbol{\Omega} := \boldsymbol{\omega}^{\wedge} = \sum_{q=1}^{3} \omega_q \mathbf{S}_q$, with $\boldsymbol{\omega} = [\omega_1\ \omega_2\ \omega_3]^\top$. Let $\mathbf{u} \in \mathbb{R}^4$ denote the control input vector, defined as $\mathbf{u} := [u_1, u_2, u_3, u_4]^\top = [f, \boldsymbol{\tau}^\top]^\top$. The state vector $\mathbf{x} \in \mathbb{R}^{18}$ is defined as $\mathbf{x} := [\mathbf{s}^\top,\ \mathbf{v}^\top,\ \underline{\mathbf{R}}^\top,\ \boldsymbol{\omega}^\top]^\top$, where $\underline{\mathbf{R}} = \mathrm{vec}(\mathbf{R})$. The equations of motion describing the *nominal rigid-body dynamics* of the quadrotor are given in a control-affine form as follows [45]

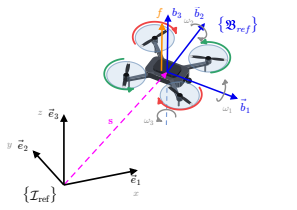


Figure 2: Quadrotor's coordinate frames.

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}, \mathbf{u}) = \mathbf{f}(\mathbf{x}) + \sum_{i=1}^{4} \mathbf{g}_i(\mathbf{x}) u_i, \text{ where} \tag{1}$$

$$\mathbf{f}(\mathbf{x}) = \left[\mathbf{v}^\top,\ -\vec{\mathbf{g}}^\top,\ (\underline{\mathbf{R}\boldsymbol{\Omega}})^\top,\ \mathbf{c}(\boldsymbol{\omega})^\top\right]^\top,\ \mathbf{g}_1(\mathbf{x}) = \left[\mathbf{0}_{1\times3},\ (\tfrac{1}{m}\mathbf{R}\vec{e}_3)^\top,\ \mathbf{0}_{1\times9},\ \mathbf{0}_{1\times3}\right]^\top, \tag{2a}$$

$$\mathbf{c}(\boldsymbol{\omega}) = -\tilde{\mathbf{J}}\boldsymbol{\Omega}\mathbf{J}\boldsymbol{\omega} = -\sum_{i=1}^{3}\sum_{j=1}^{3} \tilde{\mathbf{J}}\mathbf{S}_i\mathbf{J}\vec{e}_j\omega_i\omega_j,\ \ \mathbf{g}_{r+1}(\mathbf{x}) = [\mathbf{0}_{1\times3},\ \mathbf{0}_{1\times3},\ \mathbf{0}_{1\times9},\ (\tilde{\mathbf{J}}\vec{e}_r)^\top]^\top, \tag{2b}$$

for $r \in [1:3]$. Here, $\mathbf{F} : \mathbb{R}^{18} \times \mathbb{R}^4 \to \mathbb{R}^{18}$ is the locally Lipschitz *vector field*, $\mathbf{f} : \mathbb{R}^{18} \to \mathbb{R}^{18}$ is the *drift vector field* or simply drift, and $\mathbf{g}_i : \mathbb{R}^{18} \to \mathbb{R}^{18}$ is the *input vector field* associated with input u_i (for $i \in [1:4]$). For $q, i, j, r \in [1:3]$, the q-th component of $\mathbf{c}(\boldsymbol{\omega})$ can then be expressed as

$$[\mathbf{c}(\boldsymbol{\omega})]_q = \sum_{i=1}^{3}\sum_{j=1}^{3} \gamma_{qij}\omega_i\omega_j,\ \gamma_{qij} := -\vec{e}_q^\top \tilde{\mathbf{J}}\mathbf{S}_i\mathbf{J}\vec{e}_j, \text{ and define } \kappa_{qr} := \vec{e}_q^\top \tilde{\mathbf{J}}\vec{e}_r. \tag{3}$$

The quadrotor configuration evolves on the special Euclidean group $\mathrm{SE}(3)$ (see Definition 3), which combines its position $\mathbf{s}$ and attitude $\mathbf{R}$. Since $\boldsymbol{\Omega}$ is skew-symmetric, the kinematic equation $\dot{\mathbf{R}} = \mathbf{R}\boldsymbol{\Omega}$ preserves these rotation constraints whenever $\mathbf{R}(0) \in \mathrm{SO}(3)$. We define the physical operating envelope of the quadrotor in Definition 4.

Lie derivatives. Let $\psi : \mathbb{R}^{18} \to \mathbb{R}$ be a continuously differentiable scalar *lifting function* of the quadrotor state, whose value $\psi(\mathbf{x})$ defines a *lifted coordinate*. Its Lie derivative along the vector field $\mathbf{F}$ is defined as $\mathcal{L}_{\mathbf{F}}\psi := \nabla\psi^\top\mathbf{F} = \mathbf{F}\cdot\nabla\psi$. For a vector-valued observable $\boldsymbol{\psi} = [\psi_1, \ldots, \psi_n]^\top$, the Lie derivative is defined componentwise as $\mathcal{L}_{\mathbf{F}}\boldsymbol{\psi} := [\mathcal{L}_{\mathbf{F}}\psi_1, \ldots, \mathcal{L}_{\mathbf{F}}\psi_n]^\top$. For a control-affine system, the Lie derivative inherits the affine dependence on the inputs [33]; specifically, $\mathcal{L}_{\mathbf{F}}\psi = \mathcal{L}_{\mathbf{f}}\psi + \sum_{i=1}^{4} u_i\mathcal{L}_{\mathbf{g}_i}\psi$, where $\mathcal{L}_{\mathbf{f}}\psi := \mathbf{f}\cdot\nabla\psi$ and $\mathcal{L}_{\mathbf{g}_i}\psi := \mathbf{g}_i\cdot\nabla\psi$ denote the Lie derivatives along the drift and the i-th input vector field, respectively. Hereafter, Lie derivatives are considered along the quadrotor dynamics (1), with *drift* and *input Lie derivatives* referring to the actions of $\mathcal{L}_{\mathbf{f}}$ and $\mathcal{L}_{\mathbf{g}_i}$, for $i \in [1:4]$, respectively.

Lifted representation of nominal quadrotor dynamics. We construct a lifted-coordinate representation of the nominal quadrotor dynamics using Lie derivatives of a proposed family of geometric lifting functions. Supporting notation and preliminaries are provided in § C.3. Consider the multi-

index set $\mathbb{I}_d$ defined in (34), corresponding to degree $d \in \mathbb{N}_0$ and containing n_d elements. For any multi-index $a \in \mathbb{I}_d$, define the associated angular-velocity monomial $\mu_a(\boldsymbol{\omega}) := \omega_1^{a_1}\omega_2^{a_2}\omega_3^{a_3}$. The vector of all degree-d monomials, ordered according to the enumeration of $\mathbb{I}_d$ in (35), is defined as $\boldsymbol{\mu}_d(\boldsymbol{\omega}) := \mathrm{col}_{k\in[1:n_d]}\, \mu_{a^{(d,k)}}(\boldsymbol{\omega}) \in \mathbb{R}^{n_d}$. In particular, $\boldsymbol{\mu}_0(\boldsymbol{\omega}) = 1$ and $\boldsymbol{\mu}_1(\boldsymbol{\omega}) = \boldsymbol{\omega}$. We adopt the following geometric seed lifting functions:

$$\mathbf{p_0} := \mathbf{R}^\top \mathbf{s}, \quad \mathbb{Y}_\mathbf{0} := \mathbf{R}^\top \mathbf{v}, \quad \mathbb{Z}_\mathbf{0} := \underline{\mathbf{R}} = \mathrm{vec}(\mathbf{R}). \tag{4}$$

Building on these seed lifting functions, we introduce degree-d geometric lifting functions $\mathbf{P}_d, \mathbf{Y}_d \in \mathbb{R}^{3n_d}$ and $\mathbf{Z}_d \in \mathbb{R}^{9n_d}$, and assemble them into the degree-d lifting block $\boldsymbol{\phi}_d(\mathbf{x}) \in \mathbb{R}^{16n_d}$ as follows:

$$\mathbf{P}_d := \boldsymbol{\mu}_d \otimes \mathbf{p_0}, \quad \mathbf{Y}_d := \boldsymbol{\mu}_d \otimes \mathbb{Y}_\mathbf{0}, \quad \mathbf{Z}_d := \boldsymbol{\mu}_d \otimes \mathbb{Z}_\mathbf{0}, \quad \boldsymbol{\phi}_d(\mathbf{x}) := \mathrm{col}\,(\boldsymbol{\mu}_d, \mathbf{P}_d, \mathbf{Y}_d, \mathbf{Z}_d)\,. \tag{5}$$

Let $\boldsymbol{\varphi}$ denote the indexed family of scalar lifting functions comprising all degree-d lifting blocks:

$$\boldsymbol{\varphi} := \big([\boldsymbol{\phi}_d]_i\big)_{d\in\mathbb{N}_0,\ i\in[1:16n_d]}. \tag{6}$$

The enumeration $\boldsymbol{\varphi} = (\varphi_j)_{j=0}^{\infty}$ follows increasing degree d and the prescribed component order within each block $\boldsymbol{\phi}_d$. The computation of general expressions for the Lie derivatives of the components of $\boldsymbol{\phi}_d$, for arbitrary $d \in \mathbb{N}_0$, is presented in § C.4.

To obtain a finite-dimensional lifted representation, truncate the indexed family $\boldsymbol{\varphi}$ by retaining the components of $\boldsymbol{\phi}_d$ with $d \in [0 : K]$, where $K \in \mathbb{N}$ denotes the *truncation degree*, and define the lifted state $\widehat{\mathbf{z}} := \widehat{\boldsymbol{\Phi}}(\mathbf{x}) \in \mathbb{R}^{N_K}$, with

$$\widehat{\boldsymbol{\Phi}}(\mathbf{x}) := \mathrm{col}(\boldsymbol{\phi}_0(\mathbf{x}), \ldots, \boldsymbol{\phi}_K(\mathbf{x})) \in \mathbb{R}^{N_K}, \qquad N_K := 16\textstyle\sum_{d=0}^{K} n_d = 16\binom{K+3}{3}. \tag{7}$$

Applying the chain rule and collecting the Lie-derivative contributions associated with the retained lifting blocks yields the exact finite-dimensional identity

$$\dot{\widehat{\mathbf{z}}} = \widehat{\mathbf{A}}\widehat{\mathbf{z}} + \textstyle\sum_{i=1}^{4} u_i \widehat{\mathbf{B}}_i \widehat{\mathbf{z}} + \widehat{\mathbf{r}}_K(\mathbf{x}), \tag{8}$$

where $\widehat{\mathbf{A}}, \widehat{\mathbf{B}}_i \in \mathbb{R}^{N_K \times N_K}$ are constant matrices and the truncation residual $\widehat{\mathbf{r}}_K(\mathbf{x}) \in \mathbb{R}^{N_K}$ collects the omitted degree-$(K+1)$ terms arising in the drift Lie derivative of $\boldsymbol{\phi}_K$. The input Lie derivatives involve only retained blocks, so the residual is independent of the input. The matrix construction and residual expression are provided in § C.5.

Proposition 1 (Truncation residual bound and convergence)**.** *On the compact physical operating region $\mathcal{X}$, there exists a constant $D > 0$, independent of K and $\mathbf{x}$, such that for every $K \in \mathbb{N}$,*

$$\sup_{\mathbf{x}\in\mathcal{X}} \|\widehat{\mathbf{r}}_K(\mathbf{x})\| \le D(K+1)\sqrt{n_K}\,\hat{\omega}^{K+1}. \tag{9}$$

Consequently, the truncation residual is uniformly bounded on $\mathcal{X}$ for each fixed K. If additionally $\hat{\omega} < 1$, then $\sup_{\mathbf{x}\in\mathcal{X}} \|\widehat{\mathbf{r}}_K(\mathbf{x})\| \to 0$ as $K \to \infty$. The proof is provided in § C.6.

The first component of $\widehat{\mathbf{z}}$ is the constant lifting function $\boldsymbol{\mu}_0 = 1$ and can therefore be omitted from the dynamic state. Define $\boldsymbol{\Gamma}_K := [\mathbf{0}_{(N_K-1)\times 1}\ \mathbf{I}_{N_K-1}]$ and $\mathbf{q}_0 := \mathrm{col}(1, \mathbf{0}_{N_K-1})$. The reduced lifted state is $\mathbf{z} := \boldsymbol{\Phi}(\mathbf{x}) \in \mathbb{R}^N$, where

$$\boldsymbol{\Phi}(\mathbf{x}) := \boldsymbol{\Gamma}_K \widehat{\boldsymbol{\Phi}}(\mathbf{x}) = \mathrm{col}\,(\mathbf{P}_0, \mathbf{Y}_0, \mathbf{Z}_0, \boldsymbol{\phi}_1, \ldots, \boldsymbol{\phi}_K)\,, \qquad N := N_K - 1. \tag{10}$$

Since $\widehat{\mathbf{z}} = \mathbf{q}_0 + \boldsymbol{\Gamma}_K^\top \mathbf{z}$, the truncated dynamics (8) admit the equivalent reduced representation

$$\dot{\mathbf{z}} = \mathbf{A}\mathbf{z} + \textstyle\sum_{i=1}^{4} u_i \mathbf{B}_i \mathbf{z} + \bar{\mathbf{B}}\mathbf{u} + \mathbf{r}(\mathbf{x}), \tag{11}$$

where, for $i \in [1 : 4]$, $\mathbf{A} := \boldsymbol{\Gamma}_K \widehat{\mathbf{A}} \boldsymbol{\Gamma}_K^\top$, $\mathbf{B}_i := \boldsymbol{\Gamma}_K \widehat{\mathbf{B}}_i \boldsymbol{\Gamma}_K^\top$, $\bar{\mathbf{b}}_i := \boldsymbol{\Gamma}_K \widehat{\mathbf{B}}_i \mathbf{q}_0$, $\bar{\mathbf{B}} := [\bar{\mathbf{b}}_1\ \bar{\mathbf{b}}_2\ \bar{\mathbf{b}}_3\ \bar{\mathbf{b}}_4]$, and $\mathbf{r}(\mathbf{x}) := \boldsymbol{\Gamma}_K \widehat{\mathbf{r}}_K(\mathbf{x})$. The additive input term $\bar{\mathbf{B}}\mathbf{u}$ collects the input contributions associated with the constant observable. No constant drift term appears because $\boldsymbol{\Gamma}_K \widehat{\mathbf{A}} \mathbf{q}_0 = \mathbf{0}$. The derivation is provided in § C.7. Since the first component of $\widehat{\mathbf{r}}$ is zero, $\|\mathbf{r}(\mathbf{x})\| = \|\widehat{\mathbf{r}}(\mathbf{x})\|$; therefore, Proposition 1 also applies to $\mathbf{r}$.

Remark 1 (Physical-state reconstruction)**.** *For $K \ge 1$, define the reconstruction map $\widetilde{\boldsymbol{\Phi}} : \boldsymbol{\Phi}(\mathcal{X}) \to \mathcal{X}$ by $\widetilde{\boldsymbol{\Phi}}(\mathbf{z}) := \mathrm{col}\left(\mathrm{vec}^{-1}(\mathbf{Z}_0)\mathbf{P}_0,\ \mathrm{vec}^{-1}(\mathbf{Z}_0)\mathbf{Y}_0,\ \mathbf{Z}_0,\ \boldsymbol{\mu}_1\right)$. Then $\widetilde{\boldsymbol{\Phi}}(\boldsymbol{\Phi}(\mathbf{x})) = \mathbf{x}$ for every $\mathbf{x} \in \mathcal{X}$.*

Nominal lifted bilinear model. The omission of the truncation residual in (11) defines the nominal finite-dimensional bilinear model

$$\dot{\tilde{\mathbf{z}}} = \mathbf{A}\tilde{\mathbf{z}} + \textstyle\sum_{i=1}^{4} u_i \mathbf{B}_i \tilde{\mathbf{z}} + \bar{\mathbf{B}}\mathbf{u}, \quad \tilde{\mathbf{z}}(0) = \boldsymbol{\Phi}(\mathbf{x}(0)). \tag{12}$$

Here, $\tilde{\mathbf{z}}$ denotes the predicted lifted state. For nominal rollouts, extend the reconstruction map algebraically. Define the predicted physical state as $\tilde{\mathbf{x}}(t) := \widetilde{\boldsymbol{\Phi}}(\tilde{\mathbf{z}}(t))$. The numerical study in § C.8 evaluates open-loop prediction fidelity using identical initial physical states and input sequences. The results show that increasing the truncation degree generally improves short-horizon accuracy and

geometric consistency, with even low-degree truncations remaining close to $\mathrm{SO}(3)$ over sufficiently short horizons.

Aerodynamic resistance augmentation. We augment (1) with body-frame linear drag $\mathbf{f}^B_{\mathrm{drag}} = -\mathbf{D}_v\mathbf{R}^\top\mathbf{v}$ and angular damping $\boldsymbol{\tau}^B_{\mathrm{damp}} = -\mathbf{D}_\omega\boldsymbol{\omega},$. The constant coefficient matrices satisfy $\mathbf{D}_v = \mathbf{D}_v^\top \succeq 0$ and $\mathbf{D}_\omega = \mathbf{D}_\omega^\top \succeq 0$. Define the composite damping matrices $\boldsymbol{\Lambda}_v := m^{-1}\mathbf{D}_v$ and $\boldsymbol{\Lambda}_\omega := \tilde{\mathbf{J}}\mathbf{D}_\omega$. As derived in § C.9, the augmentation induces a structured drift correction $\Delta\mathbf{A}$. Writing $\mathbf{A}_{\mathrm{aug}} := \mathbf{A} + \Delta\mathbf{A}$, the aerodynamic resistance augmented nominal bilinear model becomes

$$\dot{\tilde{\mathbf{z}}} = \mathbf{A}_{\mathrm{aug}}\tilde{\mathbf{z}} + \textstyle\sum_{i=1}^4 u_i\mathbf{B}_i\tilde{\mathbf{z}} + \bar{\mathbf{B}}\mathbf{u}, \quad \tilde{\mathbf{z}}(0) = \boldsymbol{\Phi}(\mathbf{x}(0)). \tag{13}$$

The augmentation closes within each retained degree block, preserving the bilinear structure without introducing additional observables or truncation terms. A discrete-time form of (13) is presented in § C.10.

4 ABS-PWM

Let $\mathbf{x}_k := \mathbf{x}(t_k) \in \mathcal{X}$ denote the true quadrotor state at sampling instants $t_k := k\Delta t$, $k \in [0 : N_t]$, where $\Delta t > 0$ and $T = N_t\Delta t$ with $N_t \in \mathbb{N}$. We consider the discrete-time state-space model

$$\mathbf{x}_{k+1} = \mathscr{F}(\mathbf{x}_k, \mathbf{u}_k), \qquad \mathbf{y}_k = \mathbf{x}_k \boxplus \boldsymbol{\nu}_k, \qquad \boldsymbol{\nu}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\nu). \tag{14}$$

Here, $\mathscr{F}$ is the unknown transition map; which may capture aerodynamic, actuation, environmental, and other unmodeled effects. Sensor noise $\boldsymbol{\nu}_k \in \mathbb{R}^{12}$ has covariance $\boldsymbol{\Sigma}_\nu \in \mathbb{R}^{12\times12}$, and measurements $\mathbf{y}_k \in \widehat{\mathcal{X}} \supseteq \mathcal{X}$, where $\mathcal{X}$ is the operating envelope. The operator $\boxplus$ perturbs Euclidean state components additively and attitude multiplicatively (see § D.1). Transitions apply for $k \in [0 : N_t - 1]$ and measurements are available for $k \in [0 : N_t]$.

Although all physical state components are measured, the underlying state remains latent owing to measurement noise. Under the Markov and conditional-independence assumptions, the physical system admits a controlled hidden Markov representation [46]. We approximate this representation in the lifted coordinates using the following hidden Markov model (HMM):

$$\mathbf{z}_{k+1} = \mathbf{A}_\Delta\mathbf{z}_k + \bar{\mathbf{B}}_\Delta\mathbf{u}_k + \textstyle\sum_{i=1}^4 u_{k,i}\mathbf{B}_{\Delta i}\mathbf{z}_k + \mathbf{w}_k, \qquad \mathbf{w}_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_w), \tag{15a}$$

$$\mathbf{y}_k = \widetilde{\boldsymbol{\Phi}}(\mathbf{z}_k) + \mathbf{v}_k, \qquad \mathbf{v}_k \overset{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_v). \tag{15b}$$

Here, $\mathbf{z}_k \in \mathbb{R}^N$ is the latent state, with transition matrices discretized as in § C.10. Process noise $\mathbf{w}_k$, with covariance $\boldsymbol{\Sigma}_w \in \mathbb{R}^{N\times N}$, approximately captures truncation, discretization, and unmodeled dynamics. The observation likelihood on $\widehat{\mathcal{X}}$ uses an ambient Euclidean Gaussian approximation with effective covariance $\boldsymbol{\Sigma}_v \in \mathbb{R}^{18\times18}$, distinct from the physical noise covariance $\boldsymbol{\Sigma}_\nu \in \mathbb{R}^{12\times12}$. Conditional on $\mathbf{y}_0$, the process and observation noise sequences are mutually independent and independent of $\mathbf{z}_0$. The initial belief $\mathbf{z}_0 \mid \mathbf{y}_0 \sim \mathcal{N}(\boldsymbol{\Phi}(\mathbf{y}_0), \boldsymbol{\Sigma}_0)$, with learned $\boldsymbol{\Sigma}_0 \in \mathbb{R}^{N\times N}$, initializes the filter; measurement updates begin at $k = 1$. The HMM parameters are $\boldsymbol{\theta} := (\boldsymbol{\theta}_{\mathrm{phys}}, \boldsymbol{\Sigma}_w, \boldsymbol{\Sigma}_v, \boldsymbol{\Sigma}_0)$, where $\boldsymbol{\theta}_{\mathrm{phys}} = (m, \mathbf{J}, \mathbf{D}_v, \mathbf{D}_\omega)$ determines the transition matrices through analytical construction and discretization, while the lifting $\boldsymbol{\Phi}$ and reconstruction $\widetilde{\boldsymbol{\Phi}}$ remain fixed.

ABS-PWM components. The proposed ABS-PWM comprises an inference/representation model, a transition model, and an observation model, as described below:

$$\text{Inference}: q_{\boldsymbol{\theta}}(\mathbf{z}_k \mid \mathbf{y}_{\le k}, \mathbf{u}_{<k}) = \mathcal{N}(\boldsymbol{\eta}_{k|k}, \boldsymbol{\Sigma}_{k|k}), \tag{16a}$$

$$\text{Transition}: p_{\boldsymbol{\theta}}(\mathbf{z}_{k+1} \mid \mathbf{z}_k, \mathbf{u}_k) = \mathcal{N}\big(\mathbf{A}_\Delta\mathbf{z}_k + \bar{\mathbf{B}}_\Delta\mathbf{u}_k + \textstyle\sum_{i=1}^4 u_{k,i}\mathbf{B}_{\Delta i}\mathbf{z}_k, \boldsymbol{\Sigma}_w\big), \tag{16b}$$

$$\text{Observation}: p_{\boldsymbol{\theta}}(\mathbf{y}_k \mid \mathbf{z}_k) = \mathcal{N}\big(\widetilde{\boldsymbol{\Phi}}(\mathbf{z}_k), \boldsymbol{\Sigma}_v\big). \tag{16c}$$

Here, $q_{\boldsymbol{\theta}}$ is the Gaussian approximation to the filtering posterior induced by an extended Kalman filter (EKF) using the transition and observation models. It requires no independently learned encoder and is initialized with mean $\boldsymbol{\eta}_{0|0} = \boldsymbol{\Phi}(\mathbf{y}_0)$ and covariance $\boldsymbol{\Sigma}_{0|0} = \boldsymbol{\Sigma}_0$, with subsequent measurement updates starting at $k = 1$. The observation model uses the fixed reconstruction map and an ambient Gaussian likelihood. Together, these components provide latent-state inference, uncertainty propagation, and observation prediction.

4.1 LEARNING ABS-PWM VIA GECM

We learn $\boldsymbol{\theta}$ using a generalized expectation conditional maximization (GECM) procedure [47]–[48]. The expectation step (E-step) combines extended Kalman filtering (EKF) with Rauch–Tung–Striebel (RTS) smoothing. The conditional maximization (CM) steps improve the physical parameters numerically and optimize the covariance matrices in closed form. The analytical lifting and reconstruction maps remain fixed throughout training. Consider M independent trajectories $\mathcal{D} := \{(\mathbf{y}^{(j)}_{0:N_t}, \mathbf{u}^{(j)}_{0:N_t-1})\}_{j=1}^{M}$. For compactness, define

$$\mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) := \mathbf{A}_{\Delta}(\boldsymbol{\theta}_{\text{phys}}) + \textstyle\sum_{i=1}^{4} u^{(j)}_{k,i} \mathbf{B}_{\Delta i}(\boldsymbol{\theta}_{\text{phys}}), \quad \mathbf{b}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) := \bar{\mathbf{B}}_{\Delta}(\boldsymbol{\theta}_{\text{phys}}) \mathbf{u}^{(j)}_{k}, \quad \mathbf{h} := \widetilde{\boldsymbol{\Phi}}_{\text{ext}}. \tag{17}$$

For prescribed inputs, the transition mean is affine in the latent state. Since the initial belief is conditioned on $\mathbf{y}^{(j)}_0$, we maximize the conditional observed-data log-likelihood

$$\mathfrak{L}(\boldsymbol{\theta}) := \textstyle\sum_{j=1}^{M} \log p_{\boldsymbol{\theta}}\big(\mathbf{y}^{(j)}_{1:N_t} \mid \mathbf{y}^{(j)}_0, \mathbf{u}^{(j)}_{0:N_t-1}\big) \tag{18}$$

The initial measurement initializes the latent belief and is not counted again as an observation likelihood factor. Throughout, $r \in \mathbb{N}_0$ indexes the GECM iterations, each consisting of an E-step followed by CM-steps, as described below.

E-step. At iteration r, the parameters are fixed at $\boldsymbol{\theta}^{(r)}$. For each trajectory, the EKF is initialized with $\boldsymbol{\eta}^{(j)}_{0|0} = \boldsymbol{\Phi}(\mathbf{y}^{(j)}_0)$, $\boldsymbol{\Sigma}^{(j)}_{0|0} = \boldsymbol{\Sigma}^{(r)}_0$. The affine transition permits exact Gaussian prediction, while the EKF observation update linearizes $\mathbf{h}$ at the predicted mean. An RTS backward pass then provides a joint Gaussian smoothing distribution $q_{\text{s}}^{(r,j)}(\mathbf{z}^{(j)}_{0:N_t})$ approximating the posterior $p_{\boldsymbol{\theta}^{(r)}}\big(\mathbf{z}^{(j)}_{0:N_t} \mid \mathbf{y}^{(j)}_{0:N_t}, \mathbf{u}^{(j)}_{0:N_t-1}\big)$ [46], [49]. The smoothed means, marginal covariances, and lag-one cross-covariances are denoted by $\boldsymbol{\eta}^{(j)}_{k|N_t}$, $\boldsymbol{\Sigma}^{(j)}_{k|N_t}$, and $\boldsymbol{\Sigma}^{(j)}_{k+1,k|N_t}$, respectively. The notation $q_{\text{s}}^{(r,j)}$ distinguishes the smoothing distribution from the filtering distribution $q_{\boldsymbol{\theta}}$ in (16). The EKF and RTS recursions are provided in § D.3. Holding these smoothing distributions fixed, define the expected complete-data log-likelihood

$$\mathcal{Q}_r(\boldsymbol{\theta}) := \textstyle\sum_{j=1}^{M} \mathbb{E}_{q_{\text{s}}^{(r,j)}}\big[\log p_{\boldsymbol{\theta}}\big(\mathbf{z}^{(j)}_{0:N_t}, \mathbf{y}^{(j)}_{1:N_t} \mid \mathbf{y}^{(j)}_0, \mathbf{u}^{(j)}_{0:N_t-1}\big)\big]. \tag{19}$$

The variational decomposition [50], derived in § D.2, yields

$$\mathfrak{L}(\boldsymbol{\theta}) \geq \mathcal{Q}_r(\boldsymbol{\theta}) + \textstyle\sum_{j=1}^{M} \mathcal{H}(q_{\text{s}}^{(r,j)}), \tag{20}$$

where $\mathcal{H}(q) := -\mathbb{E}_q[\log q]$ denotes differential entropy. Expansion of $\mathcal{Q}_r$ is presented in § D.2. During the subsequent CM-steps, the smoothing distributions remain fixed, so their entropies are independent of the candidate parameters. Increasing $\mathcal{Q}_r$ therefore increases this lower bound by the same amount.

Expected residual moments. Define the transition, observation, and initial residuals as

$$\begin{aligned} \mathbf{e}^{w}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) &:= \mathbf{z}^{(j)}_{k+1} - \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) \mathbf{z}^{(j)}_{k} - \mathbf{b}_{j,k}(\boldsymbol{\theta}_{\text{phys}}), \\ \mathbf{e}^{v}_{j,k} &:= \mathbf{y}^{(j)}_{k} - \mathbf{h}(\mathbf{z}^{(j)}_{k}), \qquad \mathbf{e}^{0}_{j} := \mathbf{z}^{(j)}_{0} - \boldsymbol{\Phi}(\mathbf{y}^{(j)}_{0}). \end{aligned} \tag{21}$$

Their expected second moments are

$$\begin{aligned} \mathbf{S}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) &:= \mathbb{E}_{q_{\text{s}}^{(r,j)}}\big[\mathbf{e}^{w}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) \mathbf{e}^{w}_{j,k}(\boldsymbol{\theta}_{\text{phys}})^{\top}\big], \\ \mathbf{V}_{j,k} &:= \mathbb{E}_{q_{\text{s}}^{(r,j)}}[\mathbf{e}^{v}_{j,k} (\mathbf{e}^{v}_{j,k})^{\top}], \qquad \mathbf{S}^{0}_{j} := \mathbb{E}_{q_{\text{s}}^{(r,j)}}[\mathbf{e}^{0}_{j} (\mathbf{e}^{0}_{j})^{\top}]. \end{aligned} \tag{22}$$

The transition and initial residual moments follow directly from the smoothed means and covariances. The observation residual moment can also be evaluated analytically, despite the nonlinear reconstruction. The expressions for expected second moments are derived in § D.5.

Lemma 1 (Gaussian reconstruction and residual moments)**.** *Let* $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\eta}, \boldsymbol{\Sigma})$. *For* $\ell \in [1:18]$, *define* $\mathbf{g}_\ell := \nabla h_\ell(\boldsymbol{\eta})$ *and* $\mathbf{H}_\ell := \nabla^2 h_\ell$. *Since each component of* $\mathbf{h}$ *is at most quadratic, its mean and covariance satisfy*

$$[\bar{\mathbf{h}}]_\ell = h_\ell(\boldsymbol{\eta}) + \tfrac{1}{2} \operatorname{tr}(\mathbf{H}_\ell \boldsymbol{\Sigma}), \qquad [\boldsymbol{\Omega}_h]_{\ell s} = \mathbf{g}_\ell^{\top} \boldsymbol{\Sigma} \mathbf{g}_s + \tfrac{1}{2} \operatorname{tr}(\mathbf{H}_\ell \boldsymbol{\Sigma} \mathbf{H}_s \boldsymbol{\Sigma}). \tag{23}$$

Consequently, for a fixed observation $\mathbf{y}$,

$$\mathbb{E}[(\mathbf{y} - \mathbf{h}(\mathbf{z}))(\mathbf{y} - \mathbf{h}(\mathbf{z}))^{\top}] = (\mathbf{y} - \bar{\mathbf{h}})(\mathbf{y} - \bar{\mathbf{h}})^{\top} + \boldsymbol{\Omega}_h. \tag{24}$$

Evaluating Lemma 1 at $(\boldsymbol{\eta}^{(j)}_{k|N_t}, \boldsymbol{\Sigma}^{(j)}_{k|N_t})$ gives $\mathbf{V}_{j,k}$ exactly under the Gaussian smoothing approximation (see § D.5). These analytical moments are used in the conditional covariance update; the forward EKF retains its first-order observation linearization. The proof of Lemma 1 is provided in § D.4.

CM-step 1. With the covariance matrices fixed, we decrease

$$\mathcal{L}_r(\boldsymbol{\theta}_{\text{phys}}) := \tfrac{1}{2}\textstyle\sum_{j=1}^{M}\sum_{k=0}^{N_t-1} \operatorname{tr}\left[(\boldsymbol{\Sigma}_w^{(r)})^{-1}\mathbf{S}_{j,k}(\boldsymbol{\theta}_{\text{phys}})\right] \tag{25}$$

subject to $m > 0$, $\mathbf{J} = \mathbf{J}^\top \succ 0$, $\mathbf{D}_v = \mathbf{D}_v^\top \succeq 0$, and $\mathbf{D}_\omega = \mathbf{D}_\omega^\top \succeq 0$. Feasibility is enforced through

$$m = \exp(\alpha), \qquad \mathbf{J} = \mathbf{L}_J\mathbf{L}_J^\top, \qquad \mathbf{D}_v = \mathbf{L}_v\mathbf{L}_v^\top, \qquad \mathbf{D}_\omega = \mathbf{L}_\omega\mathbf{L}_\omega^\top,$$

where the factors are lower triangular and $\mathbf{L}_J$ has strictly positive diagonal entries. A gradient-based update with backtracking accepts only steps satisfying $\mathcal{L}_r(\boldsymbol{\theta}_{\text{phys}}^{(r+1)}) \le \mathcal{L}_r(\boldsymbol{\theta}_{\text{phys}}^{(r)})$. We use a modified Newton method with eigenvalue flooring and Armijo backtracking [51]. If no acceptable step is found, the current physical parameters are retained. The smoothing moments remain fixed throughout this optimization. Further details are provided in § D.6.

CM-step 2. We impose $\boldsymbol{\Sigma}_w \succeq \epsilon_w\mathbf{I}_N$, $\boldsymbol{\Sigma}_v \succeq \epsilon_v\mathbf{I}_{18}$, and $\boldsymbol{\Sigma}_0 \succeq \epsilon_0\mathbf{I}_N$, where the floors are fixed and strictly positive.

Lemma 2 (Constrained covariance minimizer)**.** *Let* $\mathbf{S} \succeq 0$ *have eigendecomposition* $\mathbf{S} = \mathbf{U}\operatorname{diag}(s_1,\dots,s_d)\mathbf{U}^\top$. *For* $\epsilon > 0$, *define* $\mathcal{P}_\epsilon(\mathbf{S}) := \mathbf{U}\operatorname{diag}\big(\max\{s_1,\epsilon\},\dots,\max\{s_d,\epsilon\}\big)\mathbf{U}^\top$. *Then* $\mathcal{P}_\epsilon(\mathbf{S}) = \operatorname*{argmin}_{\boldsymbol{\Sigma}\succeq\epsilon\mathbf{I}_d}\left\{\log\det\boldsymbol{\Sigma} + \operatorname{tr}(\boldsymbol{\Sigma}^{-1}\mathbf{S})\right\}$.

The proof of Lemma 2 is provided in § D.7. Applying Lemma 2 to the covariance-dependent terms of $-\mathcal{Q}_r$ gives the closed-form covariance updates

$$\begin{aligned}\boldsymbol{\Sigma}_w^{(r+1)} &= \mathcal{P}_{\epsilon_w}\left(\tfrac{1}{MN_t}\textstyle\sum_{j=1}^{M}\sum_{k=0}^{N_t-1}\mathbf{S}_{j,k}(\boldsymbol{\theta}_{\text{phys}}^{(r+1)})\right),\\ \boldsymbol{\Sigma}_v^{(r+1)} &= \mathcal{P}_{\epsilon_v}\left(\tfrac{1}{MN_t}\textstyle\sum_{j=1}^{M}\sum_{k=1}^{N_t}\mathbf{V}_{j,k}\right), \quad \boldsymbol{\Sigma}_0^{(r+1)} = \mathcal{P}_{\epsilon_0}\left(\tfrac{1}{M}\textstyle\sum_{j=1}^{M}\mathbf{S}_j^0\right).\end{aligned} \tag{26}$$

Only the process covariance update depends on the newly updated physical parameters. The initial mean remains $\boldsymbol{\Phi}(\mathbf{y}_0^{(j)})$ for each trajectory. The derivation is provided in § D.8.

Theorem 1 (Conditional surrogate improvement)**.** *Suppose* $\boldsymbol{\theta}^{(r)}$ *satisfies the physical constraints and covariance floors. For fixed EKF–RTS smoothing distributions* $\{q_{\text{s}}^{(r,j)}\}_{j=1}^{M}$, *the physical update* $\mathcal{L}_r(\boldsymbol{\theta}_{\text{phys}}^{(r+1)}) \le \mathcal{L}_r(\boldsymbol{\theta}_{\text{phys}}^{(r)})$ *followed by* (26) *preserves feasibility and satisfies* $\mathcal{Q}_r(\boldsymbol{\theta}^{(r+1)}) \ge \mathcal{Q}_r(\boldsymbol{\theta}^{(r)})$. *The fixed-distribution lower bound in* (20) *therefore does not decrease.*

The proof is provided in § D.9. Starting from a feasible $\boldsymbol{\theta}^{(0)}$, training alternates EKF–RTS inference, an accepted physical parameter update, and the closed-form covariance updates until a prescribed termination criterion is met. The complete GECM algorithm is given in Algorithm 1.

5 BUC-MPC: Bilinear Uncertainty-aware Convex MPC

BUC-MPC exploits the shared ABS-PWM structure for prediction and disturbance compensation: fixed guide trajectories yield a convex tracking program, while a wrench observer weighted by $\boldsymbol{\Sigma}_w$ and $\boldsymbol{\Sigma}_v$ estimates input-equivalent discrepancies from one-step residuals without adding MPC decision variables.

Prediction with zero-order hold. Let $\hat{\mathbf{x}}_k$ denote the estimate of the physical-state after processing the measurement $\mathbf{y}_k$ with a standard ESKF [52]. For sampling interval Δt, let $\delta = n_s\Delta t$ be the shooting interval, and N_h the number of shooting intervals, with integer $n_s \ge 1$. Here, k indexes sampling instants and l indexes shooting intervals within the prediction horizon: $l|k$ denotes a prediction made at time $k\Delta t$ for time $k\Delta t + l\delta$. At time $k\Delta t$, initialize $\bar{\mathbf{z}}_{0|k} = \boldsymbol{\Phi}(\hat{\mathbf{x}}_k)$. During prediction, the input $\mathbf{u}_{l|k} \in \mathbb{R}^4$ is held constant over $[k\Delta t + l\delta, k\Delta t + (l+1)\delta)$. Fix guide states $\mathbf{z}_{l|k}^{\text{g}}$ before solving the QP. Using the learned physical parameters $\boldsymbol{\theta}_{\text{phys}}$, construct the continuous-time operators $\mathbf{A}_{\text{aug}}$, $\bar{\mathbf{B}}$, and $\mathbf{B}_i$, $i \in [1:4]$. For each fixed guide state $\mathbf{z}_{l|k}^{\text{g}}$, define

$$\mathbf{H}(\mathbf{z}) := \bar{\mathbf{B}} + [\mathbf{B}_1\mathbf{z}\ \cdots\ \mathbf{B}_4\mathbf{z}], \quad \mathbf{A}_d := e^{\delta\mathbf{A}_{\text{aug}}}, \quad \boldsymbol{\Gamma}_\delta := \textstyle\int_0^\delta e^{s\mathbf{A}_{\text{aug}}}\,\mathrm{d}s, \ \mathbf{B}_{d,l|k} := \boldsymbol{\Gamma}_\delta\mathbf{H}(\mathbf{z}_{l|k}^{\text{g}}). \tag{27}$$

Freezing the guide yields the affine prediction

$$\bar{\mathbf{z}}_{l+1|k} = \mathbf{A}_d\bar{\mathbf{z}}_{l|k} + \mathbf{B}_{d,l|k}\mathbf{u}_{l|k}, \qquad l \in [0:N_h-1]. \tag{28}$$

§ E.2 specifies the guide construction and distinguishes this shooting model from the identification-step transition used below.

Convex tracking problem. Let $(x^{\mathrm{ref}}_{l|k}, \mathbf{u}^{\mathrm{ref}}_{l|k})$ be the reference trajectories over the MPC horizon. Let the lifted state reference be $\mathbf{z}^{\mathrm{ref}}_{l|k} := \mathbf{\Phi}(\mathbf{x}^{\mathrm{ref}}_{l|k})$. Then the optimal control problem becomes

$$\begin{aligned} \min_{\mathbf{U}_k} \quad & \textstyle\sum_{l=1}^{N_h-1} \|\bar{\mathbf{z}}_{l|k} - \mathbf{z}^{\mathrm{ref}}_{l|k}\|^2_{\mathbf{Q}} + \|\bar{\mathbf{z}}_{N_h|k} - \mathbf{z}^{\mathrm{ref}}_{N_h|k}\|^2_{\mathbf{Q}_f} + \textstyle\sum_{l=0}^{N_h-1} \|\mathbf{u}_{l|k} - \mathbf{u}^{\mathrm{ref}}_{l|k}\|^2_{\mathbf{R}_u} \\ \text{subject to} \quad & (28), \quad \bar{\mathbf{z}}_{0|k} = \mathbf{\Phi}(\hat{\mathbf{x}}_{k|k}); \quad \mathbf{u}_{\min} \le \mathbf{u}_{l|k} \le \mathbf{u}_{\max}, \quad l \in [0 : N_h - 1]; \\ & \underline{\mathbf{b}}_{l|k} \le \mathbf{C}\bar{\mathbf{z}}_{l|k} \le \overline{\mathbf{b}}_{l|k}, \quad l \in [1 : N_h - 1]; \end{aligned} \tag{29}$$

where $\mathbf{U}_k := \mathrm{col}(\mathbf{u}_{0|k}, \ldots, \mathbf{u}_{N_h-1|k})$, $\mathbf{Q} \succeq 0$, $\mathbf{Q}_f \succeq 0$, and $\mathbf{R}_u \succ 0$. The selector $\mathbf{C}$ extracts $(\mathbf{P}_0, \mathbf{Y}_0, \boldsymbol{\mu}_1)$, and the bound vectors are fixed nominal guards, constructed in § E.4. For fixed guides, condensing gives a strictly convex QP in $4N_h$ inputs, with a unique optimizer whenever feasible; § E.3 provides the proof and terminal-weight construction.

Covariance-weighted wrench observer. Define the raw measurement lift $\mathbf{z}^{\mathrm{m}}_k := \mathbf{\Phi}(\mathbf{y}_k)$. Let f_Δ denote the learned bilinear transition over Δt and $\mathbf{G}_\Delta(\mathbf{z}) := \partial f_\Delta / \partial \mathbf{u}$. For $k \ge 1$, use the previously applied command to form

$$\mathbf{e}_k := \mathbf{P}\big(\mathbf{z}^{\mathrm{m}}_k - f_\Delta(\mathbf{z}^{\mathrm{m}}_{k-1}, \mathbf{u}^{\mathrm{ap}}_{k-1})\big), \qquad \mathbf{M}_k := \mathbf{P}\mathbf{G}_\Delta(\mathbf{z}^{\mathrm{m}}_{k-1}), \tag{30}$$

where $\mathbf{P} := [\mathbf{I}_{18}\ \mathbf{0}]$ for $K \ge 1$. Locally, the residual is modeled as $\mathbf{e}_k \simeq \mathbf{M}_k \mathbf{d}_k + \boldsymbol{\epsilon}_k$, with $\mathbf{d}_k \in \mathbb{R}^4$ an additive wrench discrepancy. Using the previous stored estimate, set

$$\begin{aligned} \widehat{\mathbf{F}}_{k-1} &:= \mathbf{A}_\Delta + \textstyle\sum_{i=1}^{4} (\mathbf{u}^{\mathrm{ap}}_{k-1,i} + \widehat{\mathbf{d}}_{k-1,i}) \mathbf{B}_{\Delta i}, \quad \mathbf{J}_k := D\mathbf{\Phi}(\mathbf{y}_k), \quad \mathbf{L}_k := \mathbf{P}\mathbf{J}_k, \\ \mathbf{T}_k &:= (\mathbf{P}\widehat{\mathbf{F}}_{k-1})\mathbf{J}_{k-1}, \quad \widehat{\mathbf{S}}_k := \mathbf{P}\boldsymbol{\Sigma}_w\mathbf{P}^\top + \mathbf{L}_k\boldsymbol{\Sigma}_v\mathbf{L}_k^\top + \mathbf{T}_k\boldsymbol{\Sigma}_v\mathbf{T}_k^\top + \mathbf{E}_k, \quad \mathbf{E}_k \succ 0. \end{aligned} \tag{31}$$

The observer uses raw lifts rather than ESKF estimates, preserving the measurement-noise structure of this residual. The two observation terms account approximately for noise in the current and previous measurements. The matrix $\widehat{\mathbf{S}}_k$ is a covariance proxy: reconstruction error, noise correlations, and uncertainty in the measured regressor are omitted. § E.5 derives the exact residual identity and the additional assumptions behind this approximation.

Normalize the precision and compute the regularized estimate:

$$c_k := \max_{a \in [1:18]} [\widehat{\mathbf{S}}_k^{-1}]_{aa}, \ \mathbf{W}_k := c_k^{-1}\widehat{\mathbf{S}}_k^{-1}, \ \widetilde{\mathbf{d}}_k := (\mathbf{M}_k^\top \mathbf{W}_k \mathbf{M}_k + \lambda \mathbf{I}_4)^{-1} \mathbf{M}_k^\top \mathbf{W}_k \mathbf{e}_k, \ \lambda > 0. \tag{32}$$

For fixed $\mathbf{M}_k$ and $\widehat{\mathbf{S}}_k$, this is the unique minimizer of a weighted quadratic regression. Its Gaussian MAP interpretation uses prior precision $c_k\lambda\mathbf{I}_4$, as shown in § E.7. Smooth and limit the estimate, initializing $\widehat{\mathbf{d}}_0 = \mathbf{0}$: $\widehat{\mathbf{d}}_k = \mathrm{proj}_{\mathcal{D}}\big(\widehat{\mathbf{d}}_{k-1} + \alpha(\widetilde{\mathbf{d}}_k - \widehat{\mathbf{d}}_{k-1})\big)$, $0 < \alpha \le 1$, where $\mathcal{D} := \{\mathbf{d} : \|\mathbf{d}\|_2 \le d_{\max}\}$ in the specified wrench coordinates. Apply $\mathbf{u}^{\mathrm{ap}}_k = \mathrm{clip}_{[\mathbf{u}_{\min}, \mathbf{u}_{\max}]}(\mathbf{u}^\star_{0|k} - \widehat{\mathbf{d}}_k)$ for Δt, then update and replan. Thus feedback remains at the sampling rate even when $\delta > \Delta t$. The command clip enforces input limits; post-QP compensation does not automatically preserve predicted state bounds. The learned covariances influence control through the estimated wrench correction. The smoothing gain is a design parameter, and the statistical interpretation is local to the surrogate regression. Closed-loop stability, recursive feasibility, and exact offset-free tracking require additional analysis and are not asserted here.

6 Main Results

We choose DeepMind's MuJoCo Menagerie [53] high-fidelity physics-based simulation environment for validation and benchmarking of the proposed ABS-PWM and BUC-MPC. We evaluate ABS-PWM with BUC-MPC on the Skydio X2 quadrotor in MuJoCo Menagerie, benchmarking against KoopNet [35], SkyJEPA [8], KNODE [28], and FNN [29], each paired with its corresponding predictive control method. Experimental details and additional results are provided in § F.

Open-loop prediction fidelity. We evaluate 1 s open-loop predictions (100 steps at 100 Hz), matching the horizon used in prior quadrotor modeling studies [43] and relevant to agile-flight control [54]–[55]. As shown in Table 1, ABS-PWM has the lowest attitude error $\big(\frac{1}{2}\mathrm{tr}(\mathbf{I}_3 - \mathbf{R}^\top\widetilde{\mathbf{R}})\big)$ and matches KNODE's body-rate error under ID conditions, with slightly higher position and velocity errors. Under OOD conditions, it achieves the lowest velocity, attitude, and body-rate errors, while its position error remains close to KNODE's. Its OOD attitude and body-rate errors are 63% and 47% lower than KNODE's, respectively. ABS-PWM also has the shortest mean rollout time in both settings: about 1.3 ms per 1 s horizon, yielding a 58–64× speedup over KNODE.

(a) ID

	Position [m]	Velocity [m/s]	Attitude	Body-rate [rad/s]	Time [ms]
ABS-PWM	0.37±0.32	1.26±0.67	**0.08**±0.06	**0.75**±0.33	**1.27**
KNODE	**0.36**±0.12	**1.14**±0.50	0.11±0.12	0.75±0.47	81.30
FNN	0.53±0.21	2.21±0.89	0.53±0.31	2.35±1.07	81.37
KoopNet	0.91±0.35	1.66±0.63	0.18±0.11	1.23±0.56	4.27
SkyJEPA	0.54±0.19	2.21±0.79	0.48±0.25	2.21±1.00	1.96

(b) OOD

	Position [m]	Velocity [m/s]	Attitude	Body-rate [rad/s]	Time [ms]
ABS-PWM	0.81±0.59	**2.36**±1.32	**0.20**±0.14	**1.46**±0.71	**1.25**
KNODE	**0.77**±0.33	2.67±1.21	0.54±0.36	2.77±2.08	72.91
FNN	1.00±0.30	3.64±0.92	0.90±0.29	5.04±2.48	75.58
KoopNet	1.88±1.52	3.46±2.97	0.37±0.22	2.58±1.95	4.10
SkyJEPA	0.95±0.28	3.45±0.88	0.76±0.25	4.47±2.48	1.81

Table 1: Open-loop prediction errors over a 1 s horizon under ID and OOD conditions. Errors are reported as mean RMSE ± standard deviation, and time as mean rollout time.

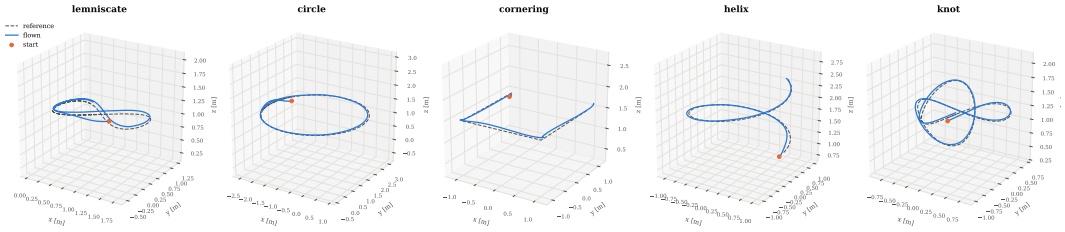


Figure 3: Representative ID trajectory-tracking demonstration with BUC-MPC for the Skydio X2.

Closed-loop trajectory tracking. We evaluate closed-loop tracking on five Skydio X2 tasks using ten runs per task, for 50 trials under each ID and OOD condition. Fig. 3 shows the tasks and representative ID trajectories tracked by BUC-MPC. As reported in Table 2, BUC-MPC completes all ID trials with the lowest mean position RMSE of 0.18 m. Under OOD conditions, it achieves the highest success rate at 86%, compared with 84% for FNN-MPC, 54% for KNODE-MPC, 18% for SkyJEPA+MPPI, and 0% for KoopNet. SkyJEPA+MPPI has a lower mean OOD RMSE of 0.66 m, but this statistic is computed from only 18% completed flights, compared with 86% for BUC-MPC. As seen in Fig. 4, BUC-MPC requires 5.24 ms per control step on average. Under the reported execution settings, it is approximately 2.9×, 1.4×, 56.6×, and 15.8× faster than KoopNet, SkyJEPA+MPPI, KNODE-MPC, and FNN-MPC, respectively. BUC-MPC remains within the 10 ms budget. Results for Crazyflie 2 are presented briefly in § F.5.

Table 2: Closed-loop performance for Skydio X2. Success rates use 50 trials per condition.

	ID		OOD	
Method	RMSE [m]	Success [%]	RMSE [m]	Success [%]
BUC-MPC	0.18±0.18	100	1.06±1.40	86
KoopNet	–	0	–	0
SkyJEPA+MPPI	0.30±0.16	42	0.66±0.33	18
KNODE-MPC	1.27±1.56	82	1.76±1.83	54
FNN-MPC	0.62±0.68	100	1.50±1.41	84

Note: Position RMSEs are reported as mean ± sample standard deviation over successful runs only.

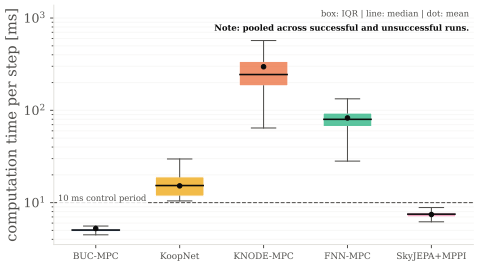


Figure 4: Per-step controller computation times.

7 Conclusion

We introduced ABS-PWM and BUC-MPC, connecting probabilistic world-model learning and efficient predictive control through a shared, physically structured bilinear representation. Analytical lifting and reconstruction provide an interpretable connection between the learned representation and the physical state, while the stochastic transition represents model and observation uncertainty. The same structure enables BUC-MPC, a convex uncertainty-aware predictive control scheme with covariance-weighted wrench compensation. Using only 90 s of noisy Skydio X2 flight data in MuJoCo environment, the framework achieves competitive open-loop prediction, accurate closed-loop tracking, and the highest OOD completion rate among the evaluated methods, with a mean control-step time of 5.24 ms. The ablations in § G demonstrate the effectiveness of the structured representation, while separately trained Crazyflie 2 experiments extend the evaluation to a second quadrotor platform. Future work will validate the framework on physical vehicles.

AI USE STATEMENT

We declare that no large language models were involved in the ideation and methodology development of this research. Large language model usage was restricted solely to GPT-6, which was utilized for polishing grammar and generating occasional data visualizations. The authors retain full responsibility for the integrity and content of the final publication.

ETHICS STATEMENT

We identified no specific ethical concerns arising from this work, which involves neither human subjects nor personal data. Physical deployment should be preceded by appropriate validation and conducted in compliance with applicable safety, privacy, and aviation regulations.

REPRODUCIBILITY STATEMENT

See § F for the full experimental setup, § B for computing resources, and § D for complete proofs of the stated propositions. We are sharing working code with this submission.

REFERENCES

[1] Nived Chebrolu, Thomas Läbe, and Cyrill Stachniss. Robust long-term registration of UAV images of crop fields for precision agriculture. *IEEE Robot. Autom. Lett.*, 3(4):3097–3104, 2018.

[2] Syed Agha Hassnain Mohsan, Nawaf Qasem Hamood Othman, Yanlong Li, Mohammed H Alsharif, and Muhammad Asghar Khan. Unmanned aerial vehicles (UAVs): Practical aspects, applications, open challenges, security issues, and future trends. *Intelligent Service Robotics*, 16(1):109–137, 2023.

[3] Haojia Li, Haokun Wang, Chen Feng, Fei Gao, Boyu Zhou, and Shaojie Shen. AutoTrans: A Complete Planning and Control Framework for Autonomous UAV Payload Transportation. *IEEE Robot. Autom. Lett.*, 8(10):6859–6866, 2023.

[4] Yiming Luo, Zixuan Zhuang, Neng Pan, Chen Feng, Shaojie Shen, Fei Gao, Hui Cheng, and Boyu Zhou. STAR-Searcher: A Complete and Efficient Aerial System for Autonomous Target Search in Complex Unknown Environments. *IEEE Robot. Autom. Lett.*, 9(5):4329–4336, 2024.

[5] Bara J Emran and Homayoun Najjaran. A review of quadrotor: An underactuated mechanical system. *Annual Reviews in Control*, 46:165–180, 2018.

[6] Bárbara Barros Carlos, Tommaso Sartor, Andrea Zanelli, Gianluca Frison, Wolfram Burgard, Moritz Diehl, and Giuseppe Oriolo. An efficient real-time NMPC for quadrotor position control under communication time delay. In *Proc. Int. Conf. Control, Autom., Robot. Vision (ICARCV)*, pages 982–989, 2020.

[7] Melissa Greeff and Angela P. Schoellig. Flatness-based model predictive control for quadrotor trajectory tracking. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, pages 6740–6745, 2018.

[8] Pratyaksh Rao, Wancong Zhang, Randall Balestriero, Yann LeCun, and Giuseppe Loianno. Skyjepa: Learning long-horizon world models for zero-shot sim-to-real control of quadrotors. *arXiv preprint arXiv:2606.23444*, 2026.

[9] A. R. P. Andriën, Erjen Lefeber, D. J. Antunes, and W. P. M. H. Heemels. Model predictive control for quadcopters with almost global trajectory tracking guarantees. *IEEE Trans. Autom. Control*, 69(8):5216–5230, 2024.

[10] Gabriel M. Hoffmann, Haomiao Huang, Steven L. Waslander, and Claire J. Tomlin. Quadrotor helicopter flight dynamics and control: Theory and experiment. In *AIAA Guidance, Navigation and Control Conference and Exhibit*, 2007. AIAA Paper 2007-6461.

[11] Leonard Bauersfeld, Elia Kaufmann, Philipp Foehn, Sihao Sun, and Davide Scaramuzza. NeuroBEM: Hybrid aerodynamic quadrotor model. In *Proc. Robot.: Sci. Syst. (RSS)*, 2021.

[12] Drew Hanover, Philipp Foehn, Sihao Sun, Elia Kaufmann, and Davide Scaramuzza. Performance, precision, and payloads: Adaptive nonlinear mpc for quadrotors. *IEEE Robotics and Automation Letters*, 7(2):690–697, 2022.

[13] Michael O'Connell, Guanya Shi, Xichen Shi, Kamyar Azizzadenesheli, Anima Anandkumar, Yisong Yue, and Soon-Jo Chung. Neural-Fly enables rapid learning for agile flight in strong winds. *Sci. Robot.*, 7(66):eabm6597, 2022.

[14] N. A. Hansen, H. Su, and X. Wang. Temporal difference learning for model predictive control. In *Proc. Int. Conf. Mach. Learn. (ICML)*, pages 8387–8406, 2022.

[15] T. Z. Jiahao, K. Y. Chee, and M. A. Hsieh. Online dynamics learning for predictive control with an application to aerial robots. In *Proc. Conf. Robot Learn. (CoRL)*, volume 205, pages 2251–2261, 2023.

[16] Alessandro Saviolo, Jonathan Frey, Abhishek Rathod, Moritz Diehl, and Giuseppe Loianno. Active learning of discrete-time dynamics for uncertainty-aware model predictive control. *IEEE Trans. Robot.*, 40:1273–1291, 2024.

[17] K. Chua, R. Calandra, R. McAllister, and S. Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Adv. Neural Inf. Process. Syst.*, volume 31, 2018.

[18] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In *Proc. Int. Conf. Mach. Learn. (ICML)*, pages 2555–2565, 2019.

[19] Maria Krinner, Elie Aljalbout, Angel Romero, and Davide Scaramuzza. Accelerating model-based reinforcement learning with state-space world models. *arXiv preprint arXiv:2502.20168*, 2025.

[20] Eashan Vytla, Bhavanishankar Kalavakolanu, Andrew Perrault, and Matthew McCrink. Dreaming falcon: Physics-informed model-based reinforcement learning for quadcopters. *arXiv preprint arXiv:2511.18243*, 2025.

[21] Angel Romero, Ashwin Shenai, Ismail Geles, Elie Aljalbout, and Davide Scaramuzza. Dream to fly: Model-based reinforcement learning for vision-based drone flight. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2026.

[22] Aderik Verraest, Stavrow Bahnam, Robin Ferede, Guido de Croon, and Christophe De Wagter. Skydreamer: Interpretable end-to-end vision-based drone racing with model-based reinforcement learning. *arXiv preprint arXiv:2510.14783*, 2025.

[23] Zian Liu, Andong Yang, Chunkai Yang, Ruidong An, Chao Gao, and Guyue Zhou. Airdreamer: Generalist drone navigation with world models. *arXiv preprint arXiv:2606.03252*, 2026.

[24] Xinhong Zhang, Runqing Wang, Yunfan Ren, Ding Yu, Boyu Zhou, Jian Sun, Fang Deng, Jie Chen, and Gang Wang. MAD: Mapping-aware world models for agile quadrotor flight. *arXiv preprint arXiv:2606.04534*, 2026.

[25] Xinhong Zhang, Qiyuan Zhu, Yubo Huang, Haolin Chen, Runqing Wang, Yuhao Mo, Zhongxin Chen, Yu Hu, Xinjiang Wang, Jian Sun, et al. Aeroact: Action-centered world-action models for language-conditioned quadrotor flight. *arXiv preprint arXiv:2607.14997*, 2026.

[26] Y. Zhang, R. Zhang, Y. Shang, Z. Yu, W. Wang, S. Yang, Q. Yan, C. Yan, and M. Feroskhan. Skytopia: Monocular drone navigation with action-conditioned latent world models. *arXiv preprint arXiv:2609.26007*, 2026.

[27] Baining Zhao, Rongze Tang, Mingyuan Jia, Ziyou Wang, Fanhang Man, Xin Zhang, Yu Shang, Weichen Zhang, Wei Wu, Chen Gao, et al. Airscape: An aerial generative world model with motion controllability. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 12519–12528, 2025.

[28] Kong Yao Chee, Tom Z. Jiahao, and M. Ani Hsieh. KNODE-MPC: A Knowledge-Based Data-Driven Predictive Control Framework for Aerial Robots. *IEEE Robot. Autom. Lett.*, 7(2):2819–2826, 2022.

[29] Jindou Jia, Zihan Yang, Meng Wang, Kexin Guo, Jianfei Yang, Xiang Yu, and Lei Guo. Feedback favors the generalization of neural ODEs. In *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2025.

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

[31] Lucian Cristian Iacob, Roland Tóth, and Maarten Schoukens. Koopman form of nonlinear systems with inputs. *Automatica*, 162:111525, 2024.

[32] Santosh Mohan Rajkumar, Dibyasri Barman, Kumar Vikram Singh, and Debdipta Goswami. On data-driven koopman representations of nonlinear delay differential equations. *arXiv preprint arXiv:2604.03086*, 2026.

[33] Debdipta Goswami and Derek A. Paley. Bilinearization, reachability, and optimal control of control-affine nonlinear systems: A koopman spectral approach. *IEEE Trans. Autom. Control*, 67(6):2715–2728, 2022.

[34] Santosh Mohan Rajkumar, Chengyu Yang, Yuliang Gu, Sheng Cheng, Naira Hovakimyan, and Debdipta Goswami. Real-time linear MPC for quadrotors on SE(3): An analytical Koopman-based realization. *IEEE Robot. Autom. Lett.*, 10(12):13018–13025, 2025.

[35] Carl Folkestad, Skylar X. Wei, and Joel W. Burdick. KoopNet: Joint learning of koopman bilinear models and function dictionaries with application to quadrotor trajectory tracking. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1344–1350, 2022.

[36] Brian Edward Jackson, Jeong Hun Lee, Kevin Tracy, and Zachary Manchester. Data-efficient model learning for control with jacobian-regularized dynamic-mode decomposition. In *Proc. Conf. Robot Learn. (CoRL)*, volume 205, pages 2273–2283, 2023.

[37] P. Huang, K. Zheng, and G. Fettweis. Data-driven Koopman operator-based error-state Kalman filter for enhanced state estimation of quadrotors in agile flight. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, pages 10983–10989, 2024.

[38] Matthias Faessler, Antonio Franchi, and Davide Scaramuzza. Differential flatness of quadrotor dynamics subject to rotor drag for accurate tracking of high-speed trajectories. *IEEE Robot. Autom. Lett.*, 3(2):620–626, 2018.

[39] Guillem Torrente, Elia Kaufmann, Philipp Foehn, and Davide Scaramuzza. Data-driven MPC for quadrotors. *IEEE Robot. Autom. Lett.*, 6(2):3769–3776, 2021.

[40] Tim Salzmann, Elia Kaufmann, Jon Arrizabalaga, Marco Pavone, Davide Scaramuzza, and Markus Ryll. Real-time neural-MPC: Deep learning model predictive control for quadrotors and agile robotic platforms. *IEEE Robot. Autom. Lett.*, 8(4):2397–2404, 2023.

[41] Alessandro Saviolo, Guanrui Li, and Giuseppe Loianno. Physics-inspired temporal learning of quadrotor dynamics for accurate model predictive trajectory tracking. *IEEE Robot. Autom. Lett.*, 7(4):10256–10263, 2022.

[42] Pratyaksh Prabhav Rao, Alessandro Saviolo, Tommaso Castiglione Ferrari, and Giuseppe Loianno. Learning long-horizon predictions for quadrotor dynamics. In *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, pages 12758–12765, 2024.

[43] Jan Węgrzynowski, Piotr Kicki, Grzegorz Czechmanowski, Maciej Piotr Krupka, and Krzysztof Walas. Beyond constant parameters: Hyper prediction models and HyperMPC. In *Proceedings of the 9th Conference on Robot Learning (CoRL)*, volume 305 of *Proceedings of Machine Learning Research*, pages 4885–4907, 2025.

[44] Avirup Das, Rishabh Dev Yadav, Sihao Sun, Mingfei Sun, Samuel Kaski, and Wei Pan. DroneDiffusion: Robust quadrotor dynamics learning with diffusion models. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 1604–1610, 2025.

[45] Taeyoung Lee, Melvin Leok, and N. Harris McClamroch. Geometric tracking control of a quadrotor UAV on $\mathrm{SE}(3)$. In *Proc. IEEE Conf. Decis. Control (CDC)*, pages 5420–5425, 2010.

[46] Simo Särkkä and Lennart Svensson. *Bayesian filtering and smoothing*, volume 17. Cambridge university press, 2023.

[47] Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *J. R. Stat. Soc. B*, 39(1):1–22, 1977.

[48] Xiao-Li Meng and Donald B. Rubin. Maximum likelihood estimation via the ECM algorithm: A general framework. *Biometrika*, 80(2):267–278, 1993.

[49] Zoubin Ghahramani and Sam Roweis. Learning nonlinear dynamical systems using an EM algorithm. In *Adv. Neural Inf. Process. Syst.*, volume 11, 1998.

[50] Radford M. Neal and Geoffrey E. Hinton. A view of the EM algorithm that justifies incremental, sparse, and other variants. In *Learning in Graphical Models*, pages 355–368. Springer, 1998.

[51] J Nocedal and S Wright. *Numerical Optimization*. Springer, 2006.

[52] Joan Sola. Quaternion kinematics for the error-state kalman filter. *arXiv preprint arXiv:1711.02508*, 2017.

[53] Kevin Zakka, Yuval Tassa, and MuJoCo Menagerie Contributors. MuJoCo Menagerie: A collection of high-quality simulation models for MuJoCo, 2022.

[54] Yunfan Ren, Zhiyuan Zhu, Jiaxu Xing, and Davide Scaramuzza. Learning Agile Quadrotor Flight in the Real World. In *Proceedings of Robotics: Science and Systems*, Sydney, Australia, July 2026.

[55] Rudolf Reiter, Chao Qin, Leonard Bauersfeld, and Davide Scaramuzza. Temporal cascading of planning and control for quadrotor MPC. *IEEE Transactions on Robotics*, 2026. Early Access.

[56] D. Mellinger and V. Kumar. Minimum snap trajectory generation and control for quadrotors. In *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pages 2520–2525, 2011.

[57] J. Levinson, C. Esteves, K. Chen, N. Snavely, A. Kanazawa, A. Rostamizadeh, and A. Makadia. An analysis of SVD for deep rotation estimation. In *Adv. Neural Inf. Process. Syst.*, volume 33, 2020.

[58] Herbert E. Rauch, Fong Tung, and Charlotte T. Striebel. Maximum likelihood estimates of linear dynamic systems. *AIAA J.*, 3(8):1445–1450, 1965.

[59] Daniel Arnström, Alberto Bemporad, and Daniel Axehill. A dual active-set solver for embedded quadratic programming using recursive LDL^T updates. *IEEE Transactions on Automatic Control*, 67(8):4362–4369, 2022.

APPENDIX CONTENTS

A Notations, Facts, Acronyms, and Basic Definitions **16**

B Computing Resources **16**

C Supplementary Material for § 3 **17**
C.1 Nominal Dynamics and Physical Operating Envelope . . . 17
C.2 Lie Derivative Example . . . 17
C.3 Notations and Preliminaries of Lifting . . . 17
C.4 Computation of Lie Derivatives . . . 18
C.5 Finite-dimensional lifted dynamics . . . 18
C.6 Proof of Proposition 1 . . . 19
C.7 Removal of the constant observable . . . 19
C.8 Prediction Fidelity of Nominal Bilinear Form . . . 20
C.9 Aerodynamic Resistance Augmentation in Lifted Representation . . . 22
C.10 Discrete-time form of Augmented Nominal Bilinear Dynamics . . . 24

D Supplementary Material for § 4 **25**
D.1 Preliminaries . . . 25
D.2 Conditional Complete-Data Objective . . . 25
D.3 EKF Filtering and RTS Smoothing . . . 26
D.4 Proof of Lemma 1 . . . 27
D.5 Expected Residual Second Moments . . . 28
D.6 Feasible Physical Parameter Updates . . . 28
D.7 Proof of Lemma 2 . . . 29
D.8 CM-step 2 Derivation . . . 30
D.9 Proof of Conditional Surrogate Improvement . . . 31

E Supplementary Material for § 5 **34**
E.1 State estimation and measurement conventions . . . 34
E.2 Guide trajectories and the two prediction intervals . . . 34
E.3 Condensing, convexity, and terminal weight . . . 34
E.4 Nominal guards and precomputed uncertainty margins . . . 35
E.5 Exact residual identity and the covariance approximation . . . 35
E.6 Analytic measurement-to-lift Jacobian . . . 36
E.7 Normalized regression and its statistical interpretation . . . 36
E.8 Small-matrix implementation and execution . . . 37
E.9 Scope of the uncertainty and control claims . . . 37

F Supplementary Material for § 6 **39**
F.1 Preliminary Info for Experiments . . . 39
F.2 Training data collection . . . 40
F.3 Test data for Open-loop Evaluation . . . 41
F.4 Closed-loop Implementation Details . . . 42
F.5 Crazyflie 2 with BUC-MPC . . . 44
F.6 Additional Results . . . 44
F.7 Failure of KoopNet Benchmark in Closed-loop . . . 44

G Ablation Study **46**
G.1 Does a learned decoder help? . . . 46

G.2 How does analytical lift, decoding, and model-based inference help ? 46
G.3 How does the structured bilinear transition help? 46
G.4 Physical-parameter initialization. 47
G.5 Observation-noise ablation. 47
G.6 Nonlinear-parametric model ablation. 47

H Limitations **48**

A Notations, Facts, Acronyms, and Basic Definitions

Table 3: Acronyms used throughout the paper.

Acronym	Definition	Acronym	Definition
ABS-PWM	Analytical Bilinear Structured-Probabilistic World Model	BUC-MPC	Bilinear Uncertainty-aware Convex Model Predictive Control
CM	Conditional Maximization	COM	Center of Mass
DAQP	Dual Active-set QP Solver	DARE	Discrete-time Algebraic Riccati Equation
DFB	Differential-Flatness-Based	EKF	Extended Kalman Filter
ESKF	Error-State Kalman Filter	FNN	Feedback Neural Network
GECM	Generalized Expectation Conditional Maximization	GPU	Graphics Processing Unit
GRU	Gated Recurrent Unit	HMM	Hidden Markov Model
ID	In-Distribution	IMU	Inertial Measurement Unit
IPOPT	Interior Point OPTimizer	KL	Kullback–Leibler Divergence
MAD	Mapping-Aware Dreamer	MAP	Maximum A Posteriori Estimate
MBRL	Model-Based Reinforcement Learning	MLP	Multilayer Perceptron
MPC	Model Predictive Control	MPE	Median Prediction Error
MPPI	Model Predictive Path Integral	NLP	Nonlinear Program
ODE	Ordinary Differential Equation	OOD	Out-of-Distribution
QP	Quadratic Program	RK4	Fourth-Order Runge–Kutta
RL	Reinforcement Learning	RMSE	Root Mean Squared Error
RTS	Rauch–Tung–Striebel Smoother	SVD	Singular Value Decomposition
WM	World Model		

Notations. Let $\mathbb{N} := \{1, 2, \cdots\}$, $\mathbb{N}_0 := \{0\} \cup \mathbb{N}$, $\mathbb{R}^+ := (0, +\infty)$, and $\mathbb{R}_0^+ := [0, +\infty)$. For $a,\ b \in \mathbb{N}_0$ with $b > a$, we define $[a : b] := [a, b] \cap \mathbb{N}_0$ and $v_{a:b} := \{v_a, \cdots, v_b\}$. Let $\mathbf{I}_m$, $\mathbf{0}_m$, $\vec{1}_n$, and $\mathbf{0}_{m\times n}$ denote, respectively, the identity matrix, square zero matrix, column vector of ones, and zero matrix of the indicated dimensions. The symbol $\mathbf{0}$ denotes a row or column vector of the appropriate dimension. For matrix or vector representations, commas or spaces separate columns, while semicolons separate rows. The vectorization operator $\mathrm{vec}(\cdot)$ maps a matrix $\mathbf{B} \in \mathbb{R}^{m\times m}$ to $\underline{\mathbf{B}} \in \mathbb{R}^{m^2}$ by stacking its columns, while $\mathrm{vec}^{-1}(\cdot)$ recovers $\mathbf{B}$ from its vectorized form. For a matrix or vector, commas or spaces separate columns, whereas semicolons separate rows. For example, both $\mathbf{A} = [1, 2, 3; 4, 5, 6]$ and $\mathbf{A} = [1\ 2\ 3; 4\ 5\ 6]$ denote $\mathbf{A} = \begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \end{bmatrix}$. The hat map $(\cdot)^\wedge : \mathbb{R}^3 \to \mathfrak{so}(3)$ associates each vector $\mathbf{a} = [a_1, a_2, a_3]^\top$ with the skew-symmetric matrix $\mathbf{a}^\wedge = [0,\ -a_3,\ a_2;\ a_3,\ 0,\ -a_1;\ -a_2,\ a_1,\ 0]$ such that $\mathbf{a}^\wedge\mathbf{b} = \mathbf{a} \times \mathbf{b}$ for every $\mathbf{b} \in \mathbb{R}^3$. Conversely, the vee map $(\cdot)^\vee : \mathfrak{so}(3) \to \mathbb{R}^3$ is the inverse of the hat map and satisfies $(\mathbf{a}^\wedge)^\vee = \mathbf{a}$ and $(\mathbf{A}^\vee)^\wedge = \mathbf{A}$ for every $\mathbf{A} \in \mathfrak{so}(3)$. The Kronecker product is denoted by $\otimes$, the gradient operator by ∇, and function composition by $\circ$. For any logical proposition P, let $\mathbb{1}(\mathsf{P})$ denote the selector function that equals 1 if P is true and 0 otherwise.

Definition 1 (Inertial and Body Fixed Frames)**.** *An inertial reference frame represents the fixed world coordinate system with respect to which the position and orientation of the quadrotor are described. Mathematically, we define the inertial frame as* $\mathcal{I}_{ref} \triangleq \{\vec{e_1}, \vec{e_2}, \vec{e_3}\}$*, where* $\vec{e_1} = [1, 0, 0]^\top$*,* $\vec{e_2} = [0, 1, 0]^\top$*, and* $\vec{e_3} = [0, 0, 1]^\top$*. A body-fixed reference frame is a rigidly attached coordinate frame to the quadrotor body. Its origin is usually located at the center of mass (COM) of the quadrotor, and its axes move and rotate together with the vehicle. Mathematically, we define the body-fixed frame as* $\mathfrak{B}_{ref} \triangleq \{\vec{b}_1, \vec{b}_2, \vec{b}_3\}$*, where* $\vec{b}_1$*,* $\vec{b}_2$*, and* $\vec{b}_3$ *are orthonormal basis vectors fixed to the quadrotor body. Thus,* $\mathcal{I}_{\mathrm{ref}}$ *remains fixed in space, while* $\mathfrak{B}_{\mathrm{ref}}$ *moves and rotates with the quadrotor, as depicted in Fig. 2.*

Definition 2 (Special Orthogonal Group $\mathrm{SO}(3)$)**.** *The special orthogonal group* $SO(3)$ *is defined as* $\mathrm{SO}(3) := \{\mathbf{R} \in \mathbb{R}^{3\times3} \mid \mathbf{R}^\top\mathbf{R} = \mathbf{I}_3,\ \mathbf{R}\mathbf{R}^\top = \mathbf{I}_3,\ \det(\mathbf{R}) = 1\}$.

Definition 3 (Special Euclidean group $\mathrm{SE}(3)$)**.** *The special Euclidean group* $\mathrm{SE}(3)$ *is defined as the set of rigid-body transformations in three-dimensional space, consisting of a rotation and a translation. Specifically,* $\mathrm{SE}(3) := \{(\mathbf{R}, \mathbf{s}) \mid \mathbf{R} \in \mathrm{SO}(3),\ \mathbf{s} \in \mathbb{R}^3\}$.

B Computing Resources

All experiments were conducted on a laptop equipped with an Intel Core i7-11850H processor, 32 GB of RAM, and an NVIDIA GeForce RTX 3080 Laptop GPU with 16 GB of VRAM.

C SUPPLEMENTARY MATERIAL FOR § 3

C.1 NOMINAL DYNAMICS AND PHYSICAL OPERATING ENVELOPE

Quadrotor nominal dynamics. The control-affine quadrotor dynamics in (1) are derived from the following equations:

$$\dot{\mathbf{s}} = \mathbf{v}, \quad \dot{\mathbf{v}} = -\vec{\mathbf{g}} + \tfrac{f}{m}\mathbf{R}\vec{e}_3, \tag{33a}$$

$$\dot{\mathbf{R}} = \mathbf{R}\boldsymbol{\Omega}, \quad \dot{\boldsymbol{\omega}} = -\tilde{\mathbf{J}}\boldsymbol{\Omega}\mathbf{J}\boldsymbol{\omega} + \tilde{\mathbf{J}}\boldsymbol{\tau}. \tag{33b}$$

In (1), note that $\underline{\mathbf{R}\boldsymbol{\Omega}} = \mathrm{vec}(\mathbf{R}\boldsymbol{\Omega})$, so that $\dot{\underline{\mathbf{R}}} = \underline{\mathbf{R}\boldsymbol{\Omega}}$.

Definition 4 (Physical operating envelope)**.** *Define the extended state space* $\widehat{\mathcal{X}} := \{\mathbf{x} = \mathrm{col}(\mathbf{s}, \mathbf{v}, \mathrm{vec}(\mathbf{R}), \boldsymbol{\omega}) \in \mathbb{R}^{18} : \mathbf{s}, \mathbf{v}, \boldsymbol{\omega} \in \mathbb{R}^3,\ \mathbf{R} \in \mathrm{SO}(3)\}$. *For an arbitrary finite time horizon* $T > 0$*, let* $\hat{s}_0, \hat{v}, \hat{\omega} > 0$ *be finite bounds on the initial position and on the velocity and angular velocity over* $t \in [0, T]$*, respectively, and set* $\hat{s} := \hat{s}_0 + T\hat{v}$*. The compact quadrotor operating region is defined as* $\mathcal{X} := \{\mathbf{x} \in \widehat{\mathcal{X}} : \|\mathbf{s}\| \le \hat{s},\ \|\mathbf{v}\| \le \hat{v},\ \|\boldsymbol{\omega}\| \le \hat{\omega}\}$*, so that* $\mathcal{X} \subset \widehat{\mathcal{X}}$*. The physically admissible input set is defined as* $\mathcal{U} := \{\mathbf{u} \in \mathbb{R}^4 : |f| \le \hat{f},\ \|\boldsymbol{\tau}\| \le \hat{\tau}\}$*, where* $\hat{f}, \hat{\tau} > 0$ *are finite thrust and torque bounds, respectively. Componentwise state and input bounds can be translated into these norm bounds as detailed in § C.1.*

Physical operating envelope. The norm bounds in Definition 4 can be obtained from componentwise lower and upper bounds. Suppose $\underline{\mathbf{s}}_0 \le \mathbf{s}(0) \le \bar{\mathbf{s}}_0$, $\underline{\mathbf{v}} \le \mathbf{v}(t) \le \bar{\mathbf{v}}$, and $\underline{\boldsymbol{\omega}} \le \boldsymbol{\omega}(t) \le \bar{\boldsymbol{\omega}}$ for all $t \in [0, T]$, with inequalities interpreted componentwise. Define $\mathbf{b}_s := \max\{|\underline{\mathbf{s}}_0|, |\bar{\mathbf{s}}_0|\}$, $\mathbf{b}_v := \max\{|\underline{\mathbf{v}}|, |\bar{\mathbf{v}}|\}$, and $\mathbf{b}_\omega := \max\{|\underline{\boldsymbol{\omega}}|, |\bar{\boldsymbol{\omega}}|\}$, where absolute values and maxima are taken componentwise. One may then choose $\hat{s}_0 = \|\mathbf{b}_s\|$, $\hat{v} = \|\mathbf{b}_v\|_2$, and $\hat{\omega} = \|\mathbf{b}_\omega\|$. Since $\dot{\mathbf{s}} = \mathbf{v}$, these bounds imply $\|\mathbf{s}(t)\| \le \hat{s}_0 + T\hat{v} = \hat{s}$. Similarly, componentwise input bounds $\underline{f} \le f(t) \le \bar{f}$ and $\underline{\boldsymbol{\tau}} \le \boldsymbol{\tau}(t) \le \bar{\boldsymbol{\tau}}$ yield $\hat{f} = \max\{|\underline{f}|, |\bar{f}|\}$ and $\hat{\tau} = \|\max\{|\underline{\boldsymbol{\tau}}|, |\bar{\boldsymbol{\tau}}|\}\|$. The resulting norm constraints contain the corresponding componentwise constraint sets.

C.2 LIE DERIVATIVE EXAMPLE

Consider the planar control-affine system with state $\mathbf{x} = [x_1\ x_2]^\top$ and input vector $\mathbf{u} = [u_1\ u_2]^\top$ and dynamics

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}, \mathbf{u}) = \mathbf{f}(\mathbf{x}) + \textstyle\sum_{i=1}^2 \mathbf{g}_i(\mathbf{x}) u_i,$$

where $\mathbf{f}(\mathbf{x}) = [x_2,\ -x_1^3]^\top$, $\mathbf{g}_1(\mathbf{x}) = [x_2^3,\ 0]^\top$, and $\mathbf{g}_2(\mathbf{x}) = [0,\ x_1^3]^\top$. Consider the scalar observable $\psi(\mathbf{x}) = x_1 x_2$, whose gradient is $\nabla\psi(\mathbf{x}) = [\partial\psi/\partial x_1,\ \partial\psi/\partial x_2]^\top = [x_2,\ x_1]^\top$. The Lie derivative along the drift is $\mathcal{L}_\mathbf{f}\psi = \nabla\psi^\top \mathbf{f} = [x_2\ x_1][x_2\ -x_1^3]^\top = x_2^2 - x_1^4$. Similarly, the input Lie derivatives are $\mathcal{L}_{\mathbf{g}_1}\psi = \nabla\psi^\top \mathbf{g}_1 = [x_2\ x_1][x_2^3\ 0]^\top = x_2^4$ and $\mathcal{L}_{\mathbf{g}_2}\psi = \nabla\psi^\top \mathbf{g}_2 = [x_2\ x_1][0\ x_1^3]^\top = x_1^4$. Therefore, $\mathcal{L}_\mathbf{F}\psi = \mathcal{L}_\mathbf{f}\psi + \sum_{i=1}^2 u_i \mathcal{L}_{\mathbf{g}_i}\psi = x_2^2 - x_1^4 + u_1 x_2^4 + u_2 x_1^4$.

Consider the vector-valued observable $\boldsymbol{\psi}(\mathbf{x}) = [x_1\ x_2\ x_1 x_2]^\top$. The Lie derivative acts componentwise: for $\mathbf{v} \in \{\mathbf{f}, \mathbf{g}_1, \mathbf{g}_2\}$, we have $\mathcal{L}_\mathbf{v}\boldsymbol{\psi} = [\mathcal{L}_\mathbf{v} x_1\ \mathcal{L}_\mathbf{v} x_2\ \mathcal{L}_\mathbf{v}(x_1 x_2)]^\top$. Thus, $\mathcal{L}_\mathbf{f}\boldsymbol{\psi} = [x_2\ -x_1^3\ x_2^2 - x_1^4]^\top$, $\mathcal{L}_{\mathbf{g}_1}\boldsymbol{\psi} = [x_2^3\ 0\ x_2^4]^\top$, and $\mathcal{L}_{\mathbf{g}_2}\boldsymbol{\psi} = [0\ x_1^3\ x_1^4]^\top$. Consequently, $\mathcal{L}_\mathbf{F}\boldsymbol{\psi} = \mathcal{L}_\mathbf{f}\boldsymbol{\psi} + \sum_{i=1}^2 u_i \mathcal{L}_{\mathbf{g}_i}\boldsymbol{\psi}$. Therefore, $\mathcal{L}_\mathbf{F}\boldsymbol{\psi} = \left[x_2 + u_1 x_2^3,\ -x_1^3 + u_2 x_1^3,\ x_2^2 - x_1^4 + u_1 x_2^4 + u_2 x_1^4\right]^\top$.

C.3 NOTATIONS AND PRELIMINARIES OF LIFTING

For a multi-index $a = (a_1, a_2, a_3) \in \mathbb{N}_0^3$, define its degree as $|a| := a_1 + a_2 + a_3$. For each $d \in \mathbb{N}_0$, define the degree-d multi-index set as

$$\mathbb{I}_d := \{a \in \mathbb{N}_0^3 : |a| = d\}, \text{ with number of elements } n_d := \tbinom{d+2}{2}. \tag{34}$$

$\mathbb{I}_d := \{a \in \mathbb{N}_0^3 : |a| = d\}$. For each $i \in [1:3]$, let $\bar{a}_i \in \mathbb{N}_0^3$ denote the i-th canonical multi-index, i.e., $\bar{a}_1 = (1,0,0)$, $\bar{a}_2 = (0,1,0)$, and $\bar{a}_3 = (0,0,1)$. Thus, $|\bar{a}_i| = 1$ and $\bar{a}_i \in \mathbb{I}_1$ for every $i \in [1:3]$. For each $d \in \mathbb{N}_0$, choose and fix an enumeration

$$\mathbb{I}_d = \{a^{(d,1)}, \ldots, a^{(d,n_d)}\}, \text{ where } a^{(d,k)} = (a_1^{(d,k)}, a_2^{(d,k)}, a_3^{(d,k)}), \text{ for every } k \in [1:n_d]. \tag{35}$$

In particular, fix the degree-zero and degree-one enumerations as $a^{(0,1)} = (0,0,0)$ and $a^{(1,i)} = \bar{a}_i$ for every $i \in [1:3]$. For each $d \in \mathbb{N}_0$ and $i \in [1:3]$, define the selector matrix $\mathbf{M}_{i,d} \in \mathbb{R}^{n_d \times n_{d+1}}$ componentwise as

$$[\mathbf{M}_{i,d}]_{\ell k} := \mathbb{1}\left(a^{(d+1,k)} = a^{(d,\ell)} + \bar{a}_i\right) \tag{36}$$

where $\ell \in [1 : n_d]$, and $k \in [1 : n_{d+1}]$. For each $d \in \mathbb{N}_0$, define $\mathbf{C}_d \in \mathbb{R}^{n_d \times n_{d+1}}$ componentwise as

$$[\mathbf{C}_d]_{\ell k} := \sum_{q=1}^{3}\sum_{i=1}^{3}\sum_{j=1}^{3} a_q^{(d,\ell)} \gamma_{qij} \mathbb{1}(a^{(d+1,k)} = a^{(d,\ell)} - \bar{a}_q + \bar{a}_i + \bar{a}_j), \tag{37}$$

where $\ell \in [1 : n_d]$ and $k \in [1 : n_{d+1}]$. For each $r \in [1 : 3]$ and $d \geq 1$, define $\mathbf{T}_{r,d} \in \mathbb{R}^{n_d \times n_{d-1}}$ componentwise as

$$[\mathbf{T}_{r,d}]_{\ell k} := \sum_{q=1}^{3} a_q^{(d,\ell)} \kappa_{qr} \mathbb{1}(a^{(d-1,k)} = a^{(d,\ell)} - \bar{a}_q), \tag{38}$$

where $\ell \in [1 : n_d]$ and $k \in [1 : n_{d-1}]$.

C.4 Computation of Lie Derivatives

Define $\mathbf{G}_g := -(\mathbf{I}_3 \otimes \vec{\mathbf{g}}^\top) \in \mathbb{R}^{3\times 9}$, so that $\mathbf{G}_g \mathbb{Z}_0 = -\mathbf{R}^\top \vec{\mathbf{g}}$. The Lie derivatives of the seed lifting functions are computed as follows:

$$\mathcal{L}_{\mathbf{f}} \mathbf{p}_0 = -\boldsymbol{\Omega}\mathbf{p}_0 + \mathbb{Y}_0,\ \mathcal{L}_{\mathbf{f}} \mathbb{Y}_0 = -\boldsymbol{\Omega}\mathbb{Y}_0 + \mathbf{G}_g \mathbb{Z}_0,\ \mathcal{L}_{\mathbf{f}} \mathbb{Z}_0 = \mathrm{vec}(\mathbf{R}\boldsymbol{\Omega}),\ \mathcal{L}_{\mathbf{g}_1} \mathbf{p}_0 = \mathbf{0}, \tag{39a}$$

$$\mathcal{L}_{\mathbf{g}_1} \mathbb{Y}_0 = \tfrac{1}{m}\vec{e}_3,\ \mathcal{L}_{\mathbf{g}_1} \mathbb{Z}_0 = \mathbf{0},\ \mathcal{L}_{\mathbf{g}_{r+1}} \mathbf{p}_0 = \mathcal{L}_{\mathbf{g}_{r+1}} \mathbb{Y}_0 = \mathbf{0},\ \mathcal{L}_{\mathbf{g}_{r+1}} \mathbb{Z}_0 = \mathbf{0},\ \text{for } r \in [1 : 3]. \tag{39b}$$

Using $\boldsymbol{\Omega} = \sum_{i=1}^{3} \omega_i \mathbf{S}_i$ and $\mathbf{S}_i^\top = -\mathbf{S}_i$, we obtain

$$-\boldsymbol{\Omega} = \sum_{i=1}^{3} \omega_i \mathbf{S}_i^\top, \qquad \mathrm{vec}(\mathbf{R}\boldsymbol{\Omega}) = \sum_{i=1}^{3} (\mathbf{S}_i^\top \otimes \mathbf{I}_3)(\omega_i \mathbb{Z}_0). \tag{40}$$

The products involving ω_i are collected in the next-degree observable blocks through $\omega_i \boldsymbol{\mu}_d = \mathbf{M}_{i,d} \boldsymbol{\mu}_{d+1}$. By the definitions of $\mathbf{C}_d$ and $\mathbf{T}_{r,d}$ in § C.3, the Lie derivatives of the degree-d angular-velocity monomial vector satisfy

$$\mathcal{L}_{\mathbf{f}} \boldsymbol{\mu}_d = \mathbf{C}_d \boldsymbol{\mu}_{d+1}, \quad \mathcal{L}_{\mathbf{g}_1} \boldsymbol{\mu}_d = \mathbf{0}, \quad \mathcal{L}_{\mathbf{g}_{r+1}} \boldsymbol{\mu}_d = \begin{cases} \mathbf{0}, & d = 0, \\ \mathbf{T}_{r,d} \boldsymbol{\mu}_{d-1}, & d \geq 1, \end{cases} \tag{41}$$

for $d \in \mathbb{N}_0$ and $r \in [1 : 3]$. For each $d \in \mathbb{N}_0$, define

$$\mathbf{G}_d^{(3)} := \mathbf{C}_d \otimes \mathbf{I}_3 + \sum_{i=1}^{3} \mathbf{M}_{i,d} \otimes \mathbf{S}_i^\top, \qquad \mathbf{G}_d^{(9)} := \mathbf{C}_d \otimes \mathbf{I}_9 + \sum_{i=1}^{3} \mathbf{M}_{i,d} \otimes (\mathbf{S}_i^\top \otimes \mathbf{I}_3), \tag{42a}$$

$$\mathbf{G}_{g,d} := \mathbf{I}_{n_d} \otimes \mathbf{G}_g, \qquad \mathbf{D}_d := \mathbf{I}_{n_d} \otimes \left(\tfrac{1}{m}\vec{e}_3\right). \tag{42b}$$

Applying the product rule to $\mathbf{P}_d = \boldsymbol{\mu}_d \otimes \mathbf{p}_0$ gives

$$\mathcal{L}_{\mathbf{f}} \mathbf{P}_d = (\mathcal{L}_{\mathbf{f}} \boldsymbol{\mu}_d) \otimes \mathbf{p}_0 + \boldsymbol{\mu}_d \otimes (\mathcal{L}_{\mathbf{f}} \mathbf{p}_0) \tag{43}$$

$$= (\mathbf{C}_d \boldsymbol{\mu}_{d+1}) \otimes \mathbf{p}_0 + \sum_{i=1}^{3} (\omega_i \boldsymbol{\mu}_d) \otimes (\mathbf{S}_i^\top \mathbf{p}_0) + \mathbf{Y}_d \tag{44}$$

$$= \mathbf{G}_d^{(3)} \mathbf{P}_{d+1} + \mathbf{Y}_d. \tag{45}$$

The same calculation for $\mathbf{Y}_d$ and $\mathbf{Z}_d$ yields the drift identities

$$\mathcal{L}_{\mathbf{f}} \mathbf{P}_d = \mathbf{G}_d^{(3)} \mathbf{P}_{d+1} + \mathbf{Y}_d, \qquad \mathcal{L}_{\mathbf{f}} \mathbf{Y}_d = \mathbf{G}_d^{(3)} \mathbf{Y}_{d+1} + \mathbf{G}_{g,d} \mathbf{Z}_d, \qquad \mathcal{L}_{\mathbf{f}} \mathbf{Z}_d = \mathbf{G}_d^{(9)} \mathbf{Z}_{d+1}. \tag{46}$$

Similarly, the input Lie derivatives satisfy

$$\mathcal{L}_{\mathbf{g}_1} \mathbf{P}_d = \mathbf{0},\ \mathcal{L}_{\mathbf{g}_1} \mathbf{Y}_d = \mathbf{D}_d \boldsymbol{\mu}_d,\ \mathcal{L}_{\mathbf{g}_1} \mathbf{Z}_d = \mathbf{0},\ \mathcal{L}_{\mathbf{g}_{r+1}} \mathbf{P}_d = \begin{cases} \mathbf{0}, & d = 0, \\ (\mathbf{T}_{r,d} \otimes \mathbf{I}_3) \mathbf{P}_{d-1}, & d \geq 1, \end{cases} \tag{47a}$$

$$\mathcal{L}_{\mathbf{g}_{r+1}} \mathbf{Y}_d = \begin{cases} \mathbf{0}, & d = 0, \\ (\mathbf{T}_{r,d} \otimes \mathbf{I}_3) \mathbf{Y}_{d-1}, & d \geq 1, \end{cases} \qquad \mathcal{L}_{\mathbf{g}_{r+1}} \mathbf{Z}_d = \begin{cases} \mathbf{0}, & d = 0, \\ (\mathbf{T}_{r,d} \otimes \mathbf{I}_9) \mathbf{Z}_{d-1}, & d \geq 1, \end{cases} \tag{47b}$$

for $d \in \mathbb{N}_0$ and $r \in [1 : 3]$.

C.5 Finite-dimensional lifted dynamics

Along any differentiable trajectory of (1), the chain rule gives

$$\dot{\widehat{\mathbf{z}}} = \mathcal{L}_{\mathbf{f}} \widehat{\boldsymbol{\Phi}}(\mathbf{x}) + \sum_{i=1}^{4} u_i \mathcal{L}_{\mathbf{g}_i} \widehat{\boldsymbol{\Phi}}(\mathbf{x}), \tag{48}$$

where the Lie derivatives act componentwise. By the identities in § C.4, each degree-d block satisfies

$$\dot{\boldsymbol{\phi}}_d = \mathbf{A}_d^{(0)} \boldsymbol{\phi}_d + \mathbf{A}_d^{(+)} \boldsymbol{\phi}_{d+1} + u_1 \mathbf{B}_{1,d}^{(0)} \boldsymbol{\phi}_d + \sum_{r=1}^{3} u_{r+1} \mathbf{B}_{r+1,d}^{(-)} \boldsymbol{\phi}_{d-1}, \tag{49}$$

with the last sum omitted when $d = 0$. The constant coefficient blocks are

$$\mathbf{A}_d^{(0)} := \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I}_{3n_d} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{G}_{g,d} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \mathbf{A}_d^{(+)} := \mathrm{blkdiag}\left(\mathbf{C}_d, \mathbf{G}_d^{(3)}, \mathbf{G}_d^{(3)}, \mathbf{G}_d^{(9)}\right), \tag{50a}$$

$$\mathbf{B}_{1,d}^{(0)} := \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{D}_d & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \ \mathbf{B}_{r+1,d}^{(-)} := \mathrm{blkdiag}\left(\mathbf{T}_{r,d}, \mathbf{T}_{r,d} \otimes \mathbf{I}_3, \mathbf{T}_{r,d} \otimes \mathbf{I}_3, \mathbf{T}_{r,d} \otimes \mathbf{I}_9\right), \tag{50b}$$

where the first three blocks are defined for $d \geq 0$, and $\mathbf{B}_{r+1,d}^{(-)}$ is defined for $d \geq 1$ and $r \in [1:3]$. Stacking (49) for $d \in [0:K]$ yield the exact identity

$$\dot{\widehat{\mathbf{z}}} = \widehat{\mathbf{A}}\widehat{\mathbf{z}} + \textstyle\sum_{i=1}^{4} u_i \widehat{\mathbf{B}}_i \widehat{\mathbf{z}} + \widehat{\mathbf{r}}_K(\mathbf{x}), \tag{51}$$

where

$$\widehat{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_0^{(0)} & \mathbf{A}_0^{(+)} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_1^{(0)} & \mathbf{A}_1^{(+)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{A}_{K-1}^{(0)} & \mathbf{A}_{K-1}^{(+)} \\ \mathbf{0} & \cdots & \cdots & \mathbf{0} & \mathbf{A}_K^{(0)} \end{bmatrix}, \ \widehat{\mathbf{B}}_{r+1} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{B}_{r+1,1}^{(-)} & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_{r+1,2}^{(-)} & \ddots & \vdots \\ \vdots & \ddots & \ddots & \mathbf{0} \\ \mathbf{0} & \cdots & \mathbf{B}_{r+1,K}^{(-)} & \mathbf{0} \end{bmatrix},$$

$$\widehat{\mathbf{B}}_1 = \mathrm{blkdiag}\left(\mathbf{B}_{1,0}^{(0)}, \ldots, \mathbf{B}_{1,K}^{(0)}\right), \quad \widehat{\mathbf{r}}_K(\mathbf{x}) := \mathrm{col}\left(\mathbf{0}_{N_K - 16n_K}, \mathbf{A}_K^{(+)} \boldsymbol{\phi}_{K+1}(\mathbf{x})\right),$$

for $r \in [1:3]$. In particular, the residual is independent of the input. Equation (51) is an exact identity along the physical trajectory.

C.6 Proof of Proposition 1

For every $a \in \mathbb{I}_{K+1}$, the angular-velocity bound gives

$$|\mu_a(\boldsymbol{\omega})| = \Pi_{i=1}^{3} |\omega_i|^{a_i} \leq \hat{\omega}^{K+1}. \tag{52}$$

Since the seed lifting functions are continuous on the compact set $\mathcal{X}$, the constant

$$b := \max\left\{1, \sup_{\mathbf{x}\in\mathcal{X}} \|\mathbf{p_0}(\mathbf{x})\|_\infty, \sup_{\mathbf{x}\in\mathcal{X}} \|\mathbb{Y}_\mathbf{0}(\mathbf{x})\|_\infty, \sup_{\mathbf{x}\in\mathcal{X}} \|\mathbb{Z}_\mathbf{0}(\mathbf{x})\|_\infty\right\} \tag{53}$$

is finite. The definition of the lifting blocks therefore implies $\|\boldsymbol{\phi}_{K+1}(\mathbf{x})\|_\infty \leq b\,\hat{\omega}^{K+1}$ uniformly on $\mathcal{X}$. To bound the coefficient matrices, define $c_\gamma := \max_{q\in[1:3]} \sum_{i,j=1}^{3} |\gamma_{qij}|$. Using $\sum_{q=1}^{3} a_q^{(K,\ell)} = K$ in the definition of $\mathbf{C}_K$ yields

$$\sum_{k=1}^{n_{K+1}} \left|[\mathbf{C}_K]_{\ell k}\right| \leq \sum_{q=1}^{3} a_q^{(K,\ell)} \sum_{i,j=1}^{3} |\gamma_{qij}| \leq K c_\gamma. \tag{54}$$

Thus, $\|\mathbf{C}_K\|_\infty \leq K c_\gamma$. Each selector $\mathbf{M}_{i,K}$ has exactly one unit entry in each row, so $\|\mathbf{M}_{i,K}\|_\infty = 1$. Consequently, the definitions of $\mathbf{G}_K^{(3)}$ and $\mathbf{G}_K^{(9)}$, together with the Kronecker-product norm identity, give

$$\|\mathbf{A}_K^{(+)}\|_\infty \leq K c_\gamma + c_S \leq c_A(K+1), \tag{55}$$

where $c_S := \sum_{i=1}^{3} \|\mathbf{S}_i^\top\|_\infty$ and $c_A := c_\gamma + c_S$ are independent of K. Here, the matrix infinity norm is the maximum absolute row sum. Since $\mathbf{A}_K^{(+)} \boldsymbol{\phi}_{K+1}$ has $16n_K$ components,

$$\|\widehat{\mathbf{r}}_K(\mathbf{x})\| = \|\mathbf{A}_K^{(+)} \boldsymbol{\phi}_{K+1}(\mathbf{x})\| \leq 4\sqrt{n_K}\, \|\mathbf{A}_K^{(+)}\|_\infty \|\boldsymbol{\phi}_{K+1}(\mathbf{x})\|_\infty \leq 4bc_A(K+1)\sqrt{n_K}\,\hat{\omega}^{K+1}.$$

Taking the supremum over $\mathcal{X}$ proves the bound in (9) with $D = 4bc_A$. Finally, $n_K = \binom{K+2}{2}$ grows quadratically, whereas $\hat{\omega}^{K+1}$ decays geometrically when $\hat{\omega} < 1$, establishing uniform convergence.

C.7 Removal of the constant observable

The coordinate restriction matrix $\boldsymbol{\Gamma}_K$ and the constant vector $\mathbf{q}_0$ satisfy

$$\boldsymbol{\Gamma}_K \boldsymbol{\Gamma}_K^\top = \mathbf{I}_N, \qquad \boldsymbol{\Gamma}_K \mathbf{q}_0 = \mathbf{0}, \qquad \boldsymbol{\Gamma}_K^\top \boldsymbol{\Gamma}_K = \mathbf{I}_{N_K} - \mathbf{q}_0 \mathbf{q}_0^\top. \tag{56}$$

Since the first component of $\widehat{\mathbf{z}}$ is $\boldsymbol{\mu}_0 = 1$, we have $\mathbf{q}_0^\top \widehat{\mathbf{z}} = 1$. Thus, $\mathbf{z} = \boldsymbol{\Gamma}_K \widehat{\mathbf{z}}$ implies

$$\widehat{\mathbf{z}} = \left(\mathbf{q}_0 \mathbf{q}_0^\top + \boldsymbol{\Gamma}_K^\top \boldsymbol{\Gamma}_K\right) \widehat{\mathbf{z}} = \mathbf{q}_0 + \boldsymbol{\Gamma}_K^\top \mathbf{z}. \tag{57}$$

Because $\boldsymbol{\Gamma}_K$ is constant, $\dot{\mathbf{z}} = \boldsymbol{\Gamma}_K \dot{\widehat{\mathbf{z}}}$. Applying $\boldsymbol{\Gamma}_K$ to (51) and substituting the preceding identity gives

$$\dot{\mathbf{z}} = \boldsymbol{\Gamma}_K \widehat{\mathbf{A}} \left(\mathbf{q}_0 + \boldsymbol{\Gamma}_K^\top \mathbf{z}\right) + \sum_{i=1}^{4} u_i \boldsymbol{\Gamma}_K \widehat{\mathbf{B}}_i \left(\mathbf{q}_0 + \boldsymbol{\Gamma}_K^\top \mathbf{z}\right) + \boldsymbol{\Gamma}_K \widehat{\mathbf{r}}(\mathbf{x}) \tag{58}$$

$$= \boldsymbol{\Gamma}_K \widehat{\mathbf{A}} \mathbf{q}_0 + \boldsymbol{\Gamma}_K \widehat{\mathbf{A}} \boldsymbol{\Gamma}_K^\top \mathbf{z} + \sum_{i=1}^{4} u_i \boldsymbol{\Gamma}_K \widehat{\mathbf{B}}_i \boldsymbol{\Gamma}_K^\top \mathbf{z} + \sum_{i=1}^{4} u_i \boldsymbol{\Gamma}_K \widehat{\mathbf{B}}_i \mathbf{q}_0 + \boldsymbol{\Gamma}_K \widehat{\mathbf{r}}(\mathbf{x}). \tag{59}$$

The vector $\widehat{\mathbf{A}} \mathbf{q}_0$ is the first column of $\widehat{\mathbf{A}}$, which is zero by its block construction. Hence, the constant drift contribution vanishes. Using the definitions of $\mathbf{A}$, $\mathbf{B}_i$, $\bar{\mathbf{b}}_i$, and $\mathbf{r}$, together with

$$\sum_{i=1}^{4} u_i \bar{\mathbf{b}}_i = [\bar{\mathbf{b}}_1 \ \bar{\mathbf{b}}_2 \ \bar{\mathbf{b}}_3 \ \bar{\mathbf{b}}_4] \mathbf{u} = \bar{\mathbf{B}} \mathbf{u}, \tag{60}$$

yields (11). This reduction removes only the constant coordinate and introduces no additional approximation.

C.8 Prediction Fidelity of Nominal Bilinear Form

We assess the prediction fidelity of the nominal bilinear model (12) against the nominal nonlinear model (1), using identical initial physical states and input sequences. Both models use the parameters

$$m = 1.325\,\mathrm{kg}, \qquad g = 9.81\,\mathrm{m/s^2}, \qquad \mathbf{J} = \begin{bmatrix} 0.03665 & 0 & -0.0021 \\ 0 & 0.02541 & 0 \\ -0.0021 & 0 & 0.0605 \end{bmatrix} \mathrm{kg \cdot m^2}. \tag{61}$$

Let $\mathbf{x}^r(t) := [\mathbf{s}^r(t); \mathbf{v}^r(t); \underline{\mathbf{R}}^r(t); \boldsymbol{\omega}^r(t)]$ denote the reference state trajectory generated using a differential-flatness-based (DFB) reference generator following [56]. We consider five tracking tasks defined by the following position references:

$$\text{Hover:} \quad \mathbf{s}^r(t) = [1;\ 1;\ 1], \tag{62a}$$

$$\text{Helix:} \quad \mathbf{s}^r(t) = [\cos(\mathfrak{w}t);\ \sin(\mathfrak{w}t);\ 0.12t], \tag{62b}$$

$$\text{Circle:} \quad \mathbf{s}^r(t) = [\cos(\mathfrak{w}t);\ \sin(\mathfrak{w}t);\ 1], \tag{62c}$$

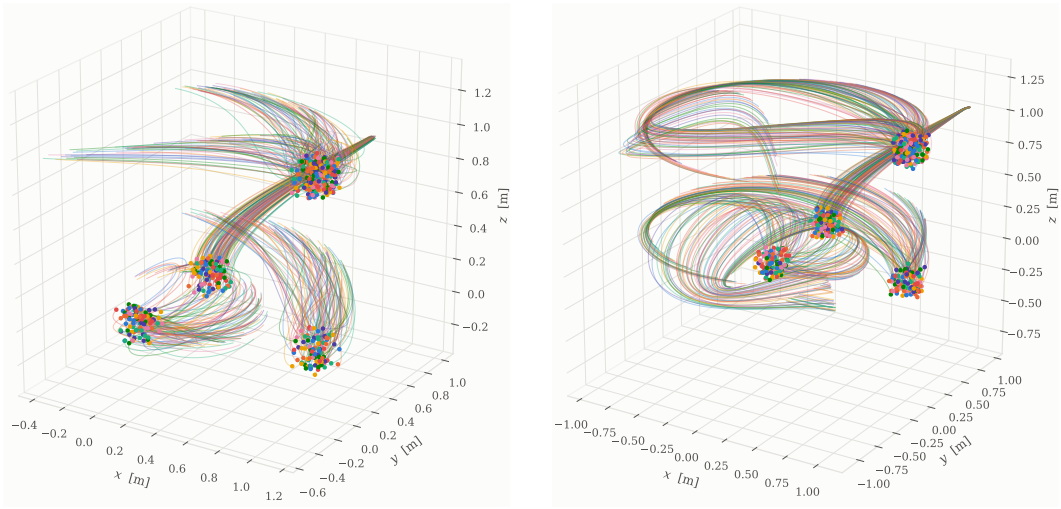
$$\text{Lemniscate:} \quad \mathbf{s}^r(t) = [\cos(\mathfrak{w}t);\ 0.5\sin(2\mathfrak{w}t);\ 1], \tag{62d}$$

$$\text{Knot:} \quad \mathbf{s}^r(t) = [0.2\sin(\mathfrak{w}t) + 0.4\sin(2\mathfrak{w}t);\ 0.2\cos(\mathfrak{w}t) - 0.4\cos(2\mathfrak{w}t);\ -0.1\sin(3\mathfrak{w}t)]. \tag{62e}$$

For each task, we generate 100 trajectories over $t \in [0, 5]$ s, sampled at 100 Hz, yielding a pooled set of 500 trajectories across the five tasks. The initial position is $\mathbf{s}(0) = [0.2;\ 0.2;\ 0.2] + \boldsymbol{\delta}$ for hover and $\mathbf{s}(0) = \mathbf{s}^r(0) + \boldsymbol{\delta}$ for all non-hover trajectories, where the components of $\boldsymbol{\delta}$ are sampled independently from $\mathrm{Unif}(-0.1, 0.1)$. All trajectories start with zero linear and angular velocities. The initial attitude is $\mathbf{R}(0) = \mathbf{I}_3$ for hover; for all non-hover trajectories, $\mathbf{R}(0)$ is constructed from Euler angles sampled independently from $\mathrm{Unif}(-5^\circ, 5^\circ)$.

For each trajectory, the input sequence is generated by applying the geometric tracking controller of Lee et al. [45] to the nominal nonlinear model, whose state trajectory $\mathbf{x}(t)$ serves as the *baseline*. The same inputs are then replayed in the nominal bilinear model, initialized with $\tilde{\mathbf{z}}(0) = \boldsymbol{\Phi}(\mathbf{x}(0))$. We assess open-loop prediction fidelity by comparing the position, velocity, rotation matrix, and angular velocity of the reconstructed prediction $\tilde{\mathbf{x}}(t) := \widetilde{\boldsymbol{\Phi}}(\tilde{\mathbf{z}}(t))$ with the corresponding components of the baseline trajectory. We consider two angular-velocity regimes, termed *moderate* ($\hat{\omega} \in (0, 1)\,\mathrm{rad/s}$) and *high* ($\hat{\omega} > 1\,\mathrm{rad/s}$) for this study. The pooled baseline position trajectories for both angular-velocity regimes are shown in Fig. 5. For non-hover trajectories, these regimes are targeted by sampling $\mathfrak{w}$ once per trajectory from $\mathrm{Unif}(0.1, 1.0)$ and $\mathrm{Unif}(0.3, 2.5)$, respectively, and holding it constant throughout the trajectory. For hover, the regimes are targeted by adjusting the geometric tracking controller gains. In all cases, the regime is verified using the resulting baseline angular velocities.

For each state component $\mathfrak{p} \in \{\mathbf{s}, \mathbf{v}, \boldsymbol{\omega}, \mathbf{R}\}$, let $\mathfrak{p}^{(j)}(t)$ and $\tilde{\mathfrak{p}}^{(j)}(t)$ denote its baseline and predicted values for trajectory $j \in [1 : 500]$, respectively. We define the median prediction error (MPE) across



(a) (b)

Figure 5: Pooled baseline position trajectories in the (a) moderate and (b) high angular-velocity regimes. Dots indicate the initial positions.

the pooled trajectories as

$$\varepsilon_{\mathfrak{p}}(t) := \mathrm{median}_{j\in[1:500]} \left\|\mathfrak{p}^{(j)}(t) - \tilde{\mathfrak{p}}^{(j)}(t)\right\|, \tag{63}$$

where $\|\cdot\|$ denotes the Euclidean norm for position, velocity, and angular velocity, and the Frobenius norm for the rotation matrix. Fig. 6a shows the MPEs of the position, velocity, angular velocity, and attitude in the moderate angular-velocity regime. Increasing the truncation degree K generally reduces the MPEs over the displayed horizon, with $K = 5$ yielding the smallest errors across all four quantities. This trend is consistent with the decay of the truncation residual established in Proposition 1. In the high angular-velocity regime (Fig. 6b), larger truncation degrees improve short-horizon accuracy, but this advantage diminishes and eventually reverses as the prediction horizon increases. By the end of the displayed horizon, higher-degree truncations exhibit substantially larger position and velocity errors, with similar reversals observed in the angular-velocity and attitude errors. These results indicate that increasing K alone does not ensure improved long-horizon prediction accuracy outside the regime covered by the truncation-residual convergence guarantee.

Consistency of predicted rotations (Geometric consistency). To quantify the departure of the predicted attitude matrices from $\mathrm{SO}(3)$, we use the SVD-based nearest-rotation projection described by Levinson et al. [57]. Let $\tilde{\mathbf{R}} = U\Sigma V^\top$, where the singular values are ordered nonincreasingly. A nearest rotation to $\tilde{\mathbf{R}}$ in the Frobenius norm is

$$\mathbf{R}^\star := U\,\mathrm{diag}\big(1, 1, \det(UV^\top)\big)\,V^\top. \tag{64}$$

For trajectory j, define the departure from $\mathrm{SO}(3)$ at time t as

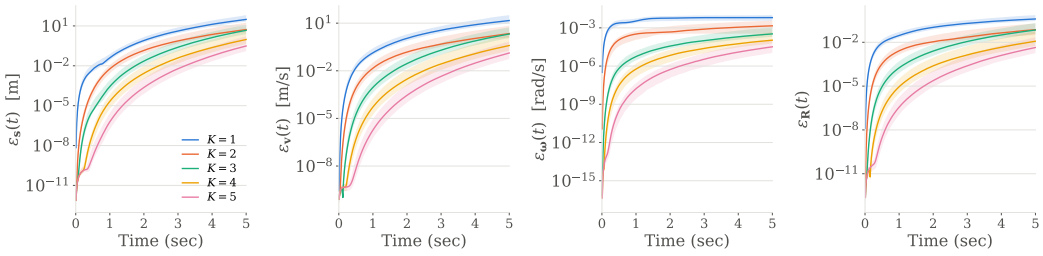
$$\varepsilon^{(j)}_{\mathrm{SO}(3)}(t) := \left\|\tilde{\mathbf{R}}^{(j)}(t) - \mathbf{R}^{\star(j)}(t)\right\|_F. \tag{65}$$

The pointwise median across the 500 pooled trajectories is

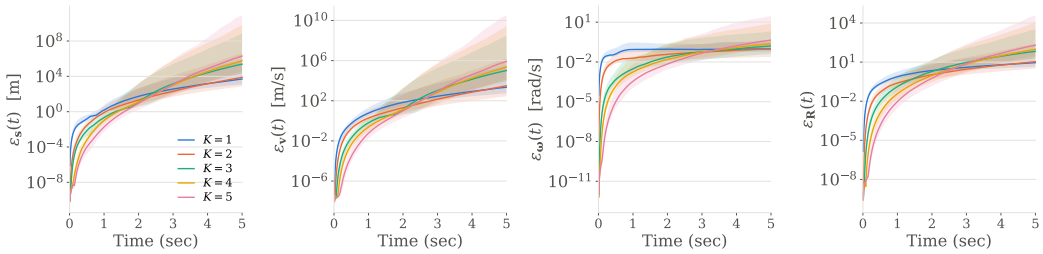
$$\bar{\varepsilon}_{\mathrm{SO}(3)}(t) := \mathrm{median}_{j\in[1:500]}\, \varepsilon^{(j)}_{\mathrm{SO}(3)}(t). \tag{66}$$

This metric measures geometric inconsistency independently of ground-truth attitude and vanishes precisely when the predicted matrix belongs to $\mathrm{SO}(3)$. The projection is used solely for evaluation and does not modify the open-loop implementation. Fig. 7 shows the median departure of the unprojected attitude predictions from $\mathrm{SO}(3)$. Even the lowest-order truncation ($K = 1$) retains approximate geometric consistency over sufficiently short horizons, while increasing K substantially reduces the short-horizon departure.

In the moderate angular-velocity regime, the median departure at $t = 1\,\mathrm{s}$ is approximately 10^{-2} for $K = 1$, a few 10^{-4} for $K = 2$, and below 10^{-6} for $K = 5$. Although departure grows with time, $K = 4$ and $K = 5$ maintain small median departures of a few 10^{-3} even at $t = 5\,\mathrm{s}$. In the high angular-velocity regime, $K = 1$ remains close to $\mathrm{SO}(3)$ over the initial portion of the rollout, but its median departure approaches unity by $t = 1\,\mathrm{s}$. At the same instant, $K = 2$ reduces the median



(a) Moderate angular-velocity regime.



(b) High angular-velocity regime.

Figure 6: Median prediction errors over time across the pooled trajectories in each angular-velocity regime for different truncation degrees K. Solid curves indicate the medians, and shaded regions denote the interquartile ranges (25th–75th percentiles).

departure to approximately 10^{-1}, while $K = 5$ yields a value below 10^{-2}. Higher truncation degrees therefore improve short-horizon geometric consistency, although this advantage eventually reverses over longer horizons, where higher-degree truncations exhibit larger departures.

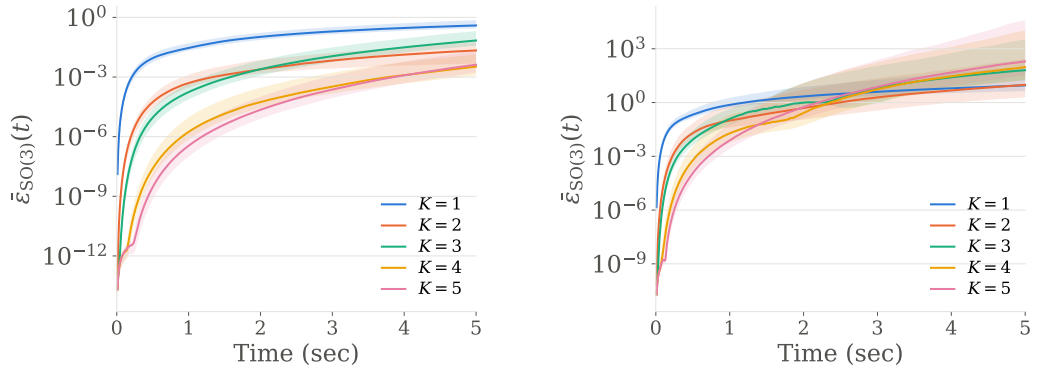


Figure 7: Departure from SO(3) : (a) moderate and (b) high angular-velocity regimes.

Together with the attitude prediction-error results, these observations support the use of unprojected predictions over suitably short horizons, including at $K = 1$ and $K = 2$, with the attainable accuracy and geometric consistency depending on both the truncation degree and the angular-velocity regimes.

C.9 Aerodynamic Resistance Augmentation in Lifted Representation

Following [38], we consider the additional body-frame force and torque representing aerodynamic resistance as

$$\mathbf{f}^B_{\text{drag}} = -\mathbf{D}_v \mathbf{R}^\top \mathbf{v}, \qquad \boldsymbol{\tau}^B_{\text{damp}} = -\mathbf{D}_\omega \boldsymbol{\omega}, \tag{67}$$

where $\mathbf{D}_v = \mathbf{D}_v^\top \succeq 0$ and $\mathbf{D}_\omega = \mathbf{D}_\omega^\top \succeq 0$. Define

$$\boldsymbol{\Lambda}_v := m^{-1}\mathbf{D}_v, \qquad \boldsymbol{\Lambda}_\omega := \tilde{\mathbf{J}}\mathbf{D}_\omega. \tag{68}$$

Using the state ordering in (1), the additional drift vector field is

$$\delta\mathbf{f}(\mathbf{x}) := \mathrm{col}\left(\mathbf{0}_{3\times1},\ -\mathbf{R}\boldsymbol{\Lambda}_v\mathbf{R}^\top\mathbf{v},\ \mathbf{0}_{9\times1},\ -\boldsymbol{\Lambda}_\omega\boldsymbol{\omega}\right), \tag{69}$$

where the zero entries denote column vectors of the indicated lengths. The augmented drift is $\mathbf{f}_{\mathrm{aug}} := \mathbf{f} + \delta\mathbf{f}$ and the aerodynamic resistance augmented dynamics become

$$\dot{\mathbf{x}} = \mathbf{f}_{\mathrm{aug}}(\mathbf{x}) + \textstyle\sum_{i=1}^{4}\mathbf{g}_i(\mathbf{x})u_i. \tag{70}$$

The augmentation leaves the input vector fields and the position and attitude kinematics unchanged. When learning the composite coefficients directly, compatibility with the assumed drag and damping matrices is enforced through

$$\boldsymbol{\Lambda}_v = \boldsymbol{\Lambda}_v^\top \succeq 0, \qquad \mathbf{J}\boldsymbol{\Lambda}_\omega = \boldsymbol{\Lambda}_\omega^\top\mathbf{J} \succeq 0. \tag{71}$$

For any continuously differentiable lifting function ψ, linearity of the Lie derivative in the vector field gives

$$\mathcal{L}_{\mathbf{f}_{\mathrm{aug}}}\psi = \mathcal{L}_{\mathbf{f}}\psi + \mathcal{L}_{\delta\mathbf{f}}\psi. \tag{72}$$

Lie derivatives of angular monomials. Use the multi-index set $\mathbb{I}_d$, its prescribed enumeration, and the canonical multi-indices $\bar{a}_i$ introduced in § C.3. For $a \in \mathbb{I}_d$,

$$\mathcal{L}_{\delta\mathbf{f}}\mu_a = -\nabla_{\boldsymbol{\omega}}\mu_a^\top\boldsymbol{\Lambda}_\omega\boldsymbol{\omega} = -\sum_{\substack{i\in[1:3]\\ a_i>0}}\sum_{j=1}^{3} a_i[\boldsymbol{\Lambda}_\omega]_{ij}\mu_{a-\bar{a}_i+\bar{a}_j}. \tag{73}$$

Define $\mathbf{E}_d(\boldsymbol{\Lambda}_\omega) \in \mathbb{R}^{n_d\times n_d}$ componentwise by

$$[\mathbf{E}_d(\boldsymbol{\Lambda}_\omega)]_{\ell k} := -\sum_{\substack{i\in[1:3]\\ a_i^{(d,\ell)}>0}}\sum_{j=1}^{3} a_i^{(d,\ell)}[\boldsymbol{\Lambda}_\omega]_{ij}\,\mathbb{1}\big(a^{(d,k)} = a^{(d,\ell)} - \bar{a}_i + \bar{a}_j\big), \tag{74}$$

for $\ell, k \in [1 : n_d]$. Then

$$\mathcal{L}_{\delta\mathbf{f}}\boldsymbol{\mu}_d = \mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\boldsymbol{\mu}_d. \tag{75}$$

In particular, $\mathbf{E}_0(\boldsymbol{\Lambda}_\omega) = \mathbf{0}$ and $\mathbf{E}_1(\boldsymbol{\Lambda}_\omega) = -\boldsymbol{\Lambda}_\omega$, since $\boldsymbol{\mu}_1 = \boldsymbol{\omega}$. Combining this expression with (41) yields

$$\mathcal{L}_{\mathbf{f}_{\mathrm{aug}}}\boldsymbol{\mu}_d = \mathbf{C}_d\boldsymbol{\mu}_{d+1} + \mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\boldsymbol{\mu}_d. \tag{76}$$

Lie derivatives of geometric lifting functions. For the seed lifting functions in (4), the additional drift contributes

$$\mathcal{L}_{\delta\mathbf{f}}\mathbf{p_0} = \mathbf{0},\ \mathcal{L}_{\delta\mathbf{f}}\mathbb{Y}_\mathbf{0} = \mathbf{R}^\top\left(-\mathbf{R}\boldsymbol{\Lambda}_v\mathbf{R}^\top\mathbf{v}\right) = -\boldsymbol{\Lambda}_v\mathbb{Y}_\mathbf{0},\ \mathcal{L}_{\delta\mathbf{f}}\mathbb{Z}_\mathbf{0} = \mathbf{0}, \tag{77}$$

where $\mathbf{R}^\top\mathbf{R} = \mathbf{I}_3$ was used. The Lie derivative satisfies the Kronecker-product rule

$$\mathcal{L}_{\delta\mathbf{f}}(\boldsymbol{\mu}_d \otimes \mathbf{p_0}) = (\mathcal{L}_{\delta\mathbf{f}}\boldsymbol{\mu}_d) \otimes \mathbf{p_0} + \boldsymbol{\mu}_d \otimes (\mathcal{L}_{\delta\mathbf{f}}\mathbf{p_0}), \tag{78}$$

and analogously for the other degree-d lifting functions. Consequently, the degree-d blocks in (5) satisfy

$$\mathcal{L}_{\delta\mathbf{f}}\mathbf{P}_d = \big(\mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_3\big)\mathbf{P}_d,\ \mathcal{L}_{\delta\mathbf{f}}\mathbf{Y}_d = \big(\mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_3 - \mathbf{I}_{n_d}\otimes\boldsymbol{\Lambda}_v\big)\mathbf{Y}_d, \tag{79a}$$

$$\mathcal{L}_{\delta\mathbf{f}}\mathbf{Z}_d = \big(\mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_9\big)\mathbf{Z}_d. \tag{79b}$$

Thus,

$$\mathcal{L}_{\delta\mathbf{f}}\boldsymbol{\phi}_d = \Delta\mathbf{A}_d\boldsymbol{\phi}_d, \tag{80}$$

where $\Delta\mathbf{A}_d \in \mathbb{R}^{16n_d\times16n_d}$ is given by

$$\Delta\mathbf{A}_d := \mathrm{blkdiag}\big(\mathbf{E}_d(\boldsymbol{\Lambda}_\omega),\ \mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_3,\ \mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_3 - \mathbf{I}_{n_d}\otimes\boldsymbol{\Lambda}_v,\ \mathbf{E}_d(\boldsymbol{\Lambda}_\omega)\otimes\mathbf{I}_9\big). \tag{81}$$

In particular,

$$\mathcal{L}_{\mathbf{f}_{\mathrm{aug}}}\boldsymbol{\phi}_d = \mathcal{L}_{\mathbf{f}}\boldsymbol{\phi}_d + \Delta\mathbf{A}_d\boldsymbol{\phi}_d. \tag{82}$$

The additional drift therefore closes within each degree-d block.

Aerodynamic resistance augmentation in finite-dimensional lifted form. For the retained degrees, define

$$\Delta\widehat{\mathbf{A}} := \mathrm{blkdiag}\big(\Delta\mathbf{A}_0,\dots,\Delta\mathbf{A}_K\big) \in \mathbb{R}^{N_K\times N_K}. \tag{83}$$

Then

$$\mathcal{L}_{\delta\mathbf{f}}\widehat{\boldsymbol{\Phi}} = \Delta\widehat{\mathbf{A}}\widehat{\boldsymbol{\Phi}}. \tag{84}$$

Using $\mathbf{\Phi} = \mathbf{\Gamma}_K \widehat{\mathbf{\Phi}}$ and $\widehat{\mathbf{\Phi}} = \mathbf{q}_0 + \mathbf{\Gamma}_K^\top \mathbf{\Phi}$, we obtain

$$\mathcal{L}_{\delta \mathbf{f}} \mathbf{\Phi} = \mathbf{\Gamma}_K \Delta \widehat{\mathbf{A}} \big(\mathbf{q}_0 + \mathbf{\Gamma}_K^\top \mathbf{\Phi} \big) \; = \Delta \mathbf{A} \mathbf{\Phi}, \quad \Delta \mathbf{A} := \mathbf{\Gamma}_K \Delta \widehat{\mathbf{A}} \mathbf{\Gamma}_K^\top \in \mathbb{R}^{N \times N}, \tag{85}$$

because $\Delta \widehat{\mathbf{A}} \mathbf{q}_0 = \mathbf{0}$. Thus, removing the constant observable introduces no constant drift term. Define

$$\mathbf{A}_{\text{aug}} := \mathbf{A} + \Delta \mathbf{A} \tag{86}$$

Thus,

$$\mathcal{L}_{\mathbf{f}_{\text{aug}}} \mathbf{\Phi}(\mathbf{x}) = \mathbf{A}_{\text{aug}} \mathbf{\Phi}(\mathbf{x}) + \mathbf{r}(\mathbf{x}). \tag{87}$$

The exact lifted state $\mathbf{z} = \mathbf{\Phi}(\mathbf{x})$ therefore satisfies the aerodynamic resistance augmented dynamics

$$\dot{\mathbf{z}} = \mathbf{A}_{\text{aug}} \mathbf{z} + \bar{\mathbf{B}} \mathbf{u} + \textstyle\sum_{i=1}^{4} \mathbf{B}_i u_i \mathbf{z} + \mathbf{r}(\mathbf{x}), \tag{88}$$

where $\mathbf{r}(\mathbf{x}; \mathbf{J}) := \mathbf{\Gamma}_K \widehat{\mathbf{r}}_K(\mathbf{x})$ is the previously defined nominal truncation residual, with its inertia dependence displayed explicitly, evaluated along the augmented physical trajectory. Since $\mathcal{L}_{\delta \mathbf{f}}$ preserves each retained degree block, the augmentation introduces no additional truncation terms. Omitting the residual gives the augmented nominal bilinear prediction model

$$\dot{\tilde{\mathbf{z}}} = \mathbf{A}_{\text{aug}} \tilde{\mathbf{z}} + \textstyle\sum_{i=1}^{4} u_i \mathbf{B}_i \tilde{\mathbf{z}} + \bar{\mathbf{B}} \mathbf{u}, \quad \tilde{\mathbf{z}}(0) = \mathbf{\Phi}(\mathbf{x}(0)). \tag{89}$$

The composite damping coefficients must satisfy (71) with the given inertia.

C.10 Discrete-time form of Augmented Nominal Bilinear Dynamics

Recalling that $t \in [0, T]$ denotes continuous time, introduce a sampling interval $\Delta t > 0$ such that $T = N_t \Delta t$ for some $N_t \in \mathbb{N}$. Define the sampling instants $t_k := k \Delta t$, $k \in \{0, \dots, N_t\}$, and write $\mathbf{z}_k := \tilde{\mathbf{z}}(t_k)$. The input is held constant over each sampling interval:

$$\mathbf{u}(t) = \mathbf{u}_k, \qquad t \in [t_k, t_{k+1}), \quad k \in \{0, \dots, N_t - 1\}.$$

Applying exponential Euler to (89), with the nominal linear drift integrated exactly and both the aerodynamic correction and input-dependent term evaluated at $\mathbf{z}_k$ and held fixed over each sampling interval, yields the discrete-time approximation

$$\mathbf{z}_{k+1} = \mathbf{A}_\Delta \mathbf{z}_k + \bar{\mathbf{B}}_\Delta \mathbf{u}_k + \textstyle\sum_{i=1}^{4} u_{k,i} \mathbf{B}_{\Delta_i} \mathbf{z}_k, \tag{90}$$

where

$$\mathbf{A}_\Delta := e^{\Delta t \mathbf{A}} + \mathbf{F}_\Delta \Delta \mathbf{A}, \qquad \bar{\mathbf{B}}_\Delta := \mathbf{F}_\Delta \bar{\mathbf{B}}, \qquad \mathbf{B}_{\Delta_i} := \mathbf{F}_\Delta \mathbf{B}_i, \qquad \mathbf{F}_\Delta := \int_0^{\Delta t} e^{s \mathbf{A}} \, ds. \tag{91}$$

Since $\mathbf{A}^{K+3} = \mathbf{0}$, the required matrix functions admit the finite expansions

$$e^{\Delta t \mathbf{A}} = \textstyle\sum_{j=0}^{K+2} \frac{(\Delta t)^j}{j!} \mathbf{A}^j, \qquad \mathbf{F}_\Delta = \textstyle\sum_{j=0}^{K+2} \frac{(\Delta t)^{j+1}}{(j+1)!} \mathbf{A}^j. \tag{92}$$

This discretization integrates the linear drift exactly and freezes the remaining input-dependent term at the beginning of each step, preserving the bilinear structure with first-order accuracy for the full dynamics.

D SUPPLEMENTARY MATERIAL FOR § 4

D.1 PRELIMINARIES

Let $\boldsymbol{\nu}_k := \mathrm{col}(\boldsymbol{\nu}_{k,p}, \boldsymbol{\nu}_{k,v}, \boldsymbol{\nu}_{k,R}, \boldsymbol{\nu}_{k,\omega}) \in \mathbb{R}^{12}$, where each component belongs to $\mathbb{R}^3$. We define

$$\mathbf{x}_k \boxplus \boldsymbol{\nu}_k := \mathrm{col}\left(\mathbf{s}_k + \boldsymbol{\nu}_{k,p},\ \mathbf{v}_k + \boldsymbol{\nu}_{k,v},\ \mathrm{vec}\big(\mathbf{R}_k \exp\big(\boldsymbol{\nu}_{k,R}^{\wedge}\big)\big),\ \boldsymbol{\omega}_k + \boldsymbol{\nu}_{k,\omega}\right),$$

where exp denotes the matrix exponential.

D.2 CONDITIONAL COMPLETE-DATA OBJECTIVE

The conditional observed-data log-likelihood in (18) is obtained by marginalizing the latent trajectories:

$$\mathfrak{L}(\boldsymbol{\theta}) = \sum_{j=1}^{M} \log \int p_{\boldsymbol{\theta}}\left(\mathbf{z}_{0:N_t}^{(j)}, \mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right) d\mathbf{z}_{0:N_t}^{(j)}. \tag{93}$$

For each trajectory, the conditional complete-data density factorizes as

$$p_{\boldsymbol{\theta}}\left(\mathbf{z}_{0:N_t}^{(j)}, \mathbf{y}_{1:N_t}^{(j)} | \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right) = p_{\boldsymbol{\theta}}(\mathbf{z}_0^{(j)} | \mathbf{y}_0^{(j)}) \prod_{k=0}^{N_t-1} p_{\boldsymbol{\theta}}(\mathbf{z}_{k+1}^{(j)} | \mathbf{z}_k^{(j)}, \mathbf{u}_k^{(j)}) \prod_{k=1}^{N_t} p_{\boldsymbol{\theta}}(\mathbf{y}_k^{(j)} | \mathbf{z}_k^{(j)}). \tag{94}$$

The initial density is Gaussian with mean $\boldsymbol{\Phi}(\mathbf{y}_0^{(j)})$ and covariance $\boldsymbol{\Sigma}_0$. The initial measurement conditions this density and is not counted again as an observation likelihood factor.

At iteration r, let $q_{\mathrm{s}}^{(r,j)}$ denote the joint Gaussian smoothing approximation constructed by the EKF–RTS recursions in § D.3. With these distributions held fixed, the expected complete-data log-likelihood is

$$\mathcal{Q}_r(\boldsymbol{\theta}) = \sum_{j=1}^{M} \mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\left[\log p_{\boldsymbol{\theta}}\left(\mathbf{z}_{0:N_t}^{(j)}, \mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right)\right].$$

Taking the logarithm of (94) and using linearity of expectation gives

$$\begin{aligned}\mathcal{Q}_r(\boldsymbol{\theta}) = \sum_{j=1}^{M} \Big\{ &\mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\big[\log p_{\boldsymbol{\theta}}(\mathbf{z}_0^{(j)} \mid \mathbf{y}_0^{(j)})\big] + \sum_{k=0}^{N_t-1} \mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\big[\log p_{\boldsymbol{\theta}}(\mathbf{z}_{k+1}^{(j)} \mid \mathbf{z}_k^{(j)}, \mathbf{u}_k^{(j)})\big] \\ &+ \sum_{k=1}^{N_t} \mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\big[\log p_{\boldsymbol{\theta}}(\mathbf{y}_k^{(j)} \mid \mathbf{z}_k^{(j)})\big]\Big\}.\end{aligned}$$

To expand these terms, recall that for a d-dimensional residual $\mathbf{e}$ and a positive definite covariance $\boldsymbol{\Sigma}$,

$$-\log \mathcal{N}(\mathbf{e}; \mathbf{0}, \boldsymbol{\Sigma}) = \tfrac{d}{2}\log(2\pi) + \tfrac{1}{2}\log\det\boldsymbol{\Sigma} + \tfrac{1}{2}\mathbf{e}^\top\boldsymbol{\Sigma}^{-1}\mathbf{e}.$$

Moreover, $\mathbb{E}_q[\mathbf{e}^\top\boldsymbol{\Sigma}^{-1}\mathbf{e}] = \mathrm{tr}\left(\boldsymbol{\Sigma}^{-1}\mathbb{E}_q[\mathbf{e}\mathbf{e}^\top]\right)$. Applying these identities to the initial, transition, and observation residuals in (21) yields

$$\begin{aligned}-\mathcal{Q}_r(\boldsymbol{\theta}) = c &+ \tfrac{M}{2}\log\det\boldsymbol{\Sigma}_0 + \tfrac{1}{2}\sum_{j=1}^{M}\mathrm{tr}(\boldsymbol{\Sigma}_0^{-1}\mathbf{S}_j^0) \\ &+ \tfrac{MN_t}{2}\log\det\boldsymbol{\Sigma}_w + \tfrac{1}{2}\sum_{j=1}^{M}\sum_{k=0}^{N_t-1}\mathrm{tr}\left(\boldsymbol{\Sigma}_w^{-1}\mathbf{S}_{j,k}(\boldsymbol{\theta}_{\mathrm{phys}})\right) \\ &+ \tfrac{MN_t}{2}\log\det\boldsymbol{\Sigma}_v + \tfrac{1}{2}\sum_{j=1}^{M}\sum_{k=1}^{N_t}\mathrm{tr}\left(\boldsymbol{\Sigma}_v^{-1}\mathbf{V}_{j,k}\right),\end{aligned} \tag{95}$$

where the expected residual outer products $\mathbf{S}_j^0$, $\mathbf{S}_{j,k}(\boldsymbol{\theta}_{\mathrm{phys}})$, and $\mathbf{V}_{j,k}$ are defined in (22). The Gaussian normalization terms contribute

$$c = \tfrac{M}{2}\big[N + N_t(N+18)\big]\log(2\pi),$$

which is independent of both $\boldsymbol{\theta}$ and r. Each trajectory contributes one N-dimensional initial residual, N_t N-dimensional transition residuals, and N_t 18-dimensional observation residuals. We express the objective as $-\mathcal{Q}_r$ so that the conditional maximization steps can be written equivalently as minimization problems. The negative sign therefore reflects only this optimization convention. Equation (95) expands the expected complete-data objective; its relation to the observed-data log-likelihood, including the entropy and KL terms, is established below.

Relation to the observed-data objective. For trajectory j, define the exact conditional smoothing density $p_{\boldsymbol{\theta}}^{(j)} := p_{\boldsymbol{\theta}}\big(\mathbf{z}_{0:N_t}^{(j)} \mid \mathbf{y}_{0:N_t}^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\big)$. For a density q, let $\mathcal{H}(q) := -\mathbb{E}_q[\log q]$ denote its differential entropy. The observed-data objective admits the decomposition[50]

$$\mathfrak{L}(\boldsymbol{\theta}) = \mathcal{Q}_r(\boldsymbol{\theta}) + \sum_{j=1}^{M} \mathcal{H}(q_{\mathrm{s}}^{(r,j)}) + \sum_{j=1}^{M} \mathrm{KL}\left(q_{\mathrm{s}}^{(r,j)} \,\|\, p_{\boldsymbol{\theta}}^{(j)}\right). \tag{96}$$

To derive this identity, apply Bayes' rule to the conditional complete-data density:

$$p_{\boldsymbol{\theta}}\left(\mathbf{z}_{0:N_t}^{(j)}, \mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right) = p_{\boldsymbol{\theta}}^{(j)}\left(\mathbf{z}_{0:N_t}^{(j)}\right) p_{\boldsymbol{\theta}}\left(\mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right).$$

Taking logarithms and expectations under $q_{\mathrm{s}}^{(r,j)}$ gives

$$\begin{aligned}
&\mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\left[\log p_{\boldsymbol{\theta}}\left(\mathbf{z}_{0:N_t}^{(j)}, \mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right)\right] \\
&= \mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\left[\log p_{\boldsymbol{\theta}}^{(j)}\left(\mathbf{z}_{0:N_t}^{(j)}\right)\right] + \log p_{\boldsymbol{\theta}}\left(\mathbf{y}_{1:N_t}^{(j)} \mid \mathbf{y}_0^{(j)}, \mathbf{u}_{0:N_t-1}^{(j)}\right).
\end{aligned}$$

By the definitions of differential entropy and KL divergence,

$$\mathbb{E}_{q_{\mathrm{s}}^{(r,j)}}\left[\log p_{\boldsymbol{\theta}}^{(j)}\left(\mathbf{z}_{0:N_t}^{(j)}\right)\right] = -\mathcal{H}(q_{\mathrm{s}}^{(r,j)}) - \mathrm{KL}\left(q_{\mathrm{s}}^{(r,j)} \,\|\, p_{\boldsymbol{\theta}}^{(j)}\right).$$

Substituting this expression and summing over $j = 1, \ldots, M$ yields

$$\mathcal{Q}_r(\boldsymbol{\theta}) = \mathfrak{L}(\boldsymbol{\theta}) - \sum_{j=1}^{M} \mathcal{H}(q_{\mathrm{s}}^{(r,j)}) - \sum_{j=1}^{M} \mathrm{KL}\left(q_{\mathrm{s}}^{(r,j)} \,\|\, p_{\boldsymbol{\theta}}^{(j)}\right).$$

Here, the expected complete-data log-densities sum to $\mathcal{Q}_r(\boldsymbol{\theta})$ and the conditional observed-data log-likelihoods sum to $\mathfrak{L}(\boldsymbol{\theta})$. Rearranging gives (96). Since each KL divergence is nonnegative,

$$\mathcal{B}_r(\boldsymbol{\theta}) := \mathcal{Q}_r(\boldsymbol{\theta}) + \sum_{j=1}^{M} \mathcal{H}(q_{\mathrm{s}}^{(r,j)}) \le \mathfrak{L}(\boldsymbol{\theta}). \tag{97}$$

During the conditional parameter updates, the smoothing distributions $q_{\mathrm{s}}^{(r,j)}$ remain fixed. Their entropies are therefore constant, so any increase in $\mathcal{Q}_r$ produces an equal increase in the lower bound $\mathcal{B}_r$. However, the exact smoothing densities $p_{\boldsymbol{\theta}}^{(j)}$ vary with the candidate parameters, and hence the KL gap generally changes as well. An increase in $\mathcal{B}_r$ therefore does not, by itself, imply an increase in the observed-data log-likelihood $\mathfrak{L}$.

D.3 EKF Filtering and RTS Smoothing

At iteration r, define

$$\mathbf{F}_{j,k}^{(r)} := \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\mathrm{phys}}^{(r)}), \qquad \mathbf{b}_{j,k}^{(r)} := \mathbf{b}_{j,k}(\boldsymbol{\theta}_{\mathrm{phys}}^{(r)}).$$

All filtering and smoothing moments in this subsection are computed using $\boldsymbol{\theta}^{(r)}$; their iteration superscripts are suppressed.

Initialization and prediction. For each trajectory, initialize

$$\boldsymbol{\eta}_{0|0}^{(j)} = \boldsymbol{\Phi}(\mathbf{y}_0^{(j)}), \qquad \boldsymbol{\Sigma}_{0|0}^{(j)} = \boldsymbol{\Sigma}_0^{(r)}.$$

For $k \in [1 : N_t]$, compute

$$\begin{aligned}
\boldsymbol{\eta}_{k|k-1}^{(j)} &= \mathbf{F}_{j,k-1}^{(r)} \boldsymbol{\eta}_{k-1|k-1}^{(j)} + \mathbf{b}_{j,k-1}^{(r)}, \\
\boldsymbol{\Sigma}_{k|k-1}^{(j)} &= \mathbf{F}_{j,k-1}^{(r)} \boldsymbol{\Sigma}_{k-1|k-1}^{(j)} (\mathbf{F}_{j,k-1}^{(r)})^{\top} + \boldsymbol{\Sigma}_w^{(r)}.
\end{aligned} \tag{98}$$

For prescribed inputs, the transition is affine in the latent state. Therefore, prediction is exact under the current Gaussian filtering approximation.

EKF observation update. Following the EKF formulation [46], we linearize the reconstruction map at the predicted mean:

$$\mathbf{h}(\mathbf{z}) \approx \mathbf{h}(\boldsymbol{\eta}_{k|k-1}^{(j)}) + \mathbf{G}_{j,k}^{-}(\mathbf{z} - \boldsymbol{\eta}_{k|k-1}^{(j)}), \qquad \mathbf{G}_{j,k}^{-} := \left.\frac{\partial \mathbf{h}}{\partial \mathbf{z}}\right|_{\boldsymbol{\eta}_{k|k-1}^{(j)}}.$$

The innovation covariance and Kalman gain are

$$\begin{aligned}\mathbf{W}_{j,k} &= \mathbf{G}^{-}_{j,k}\boldsymbol{\Sigma}^{(j)}_{k|k-1}(\mathbf{G}^{-}_{j,k})^{\top} + \boldsymbol{\Sigma}^{(r)}_{v},\\ \mathbf{K}_{j,k} &= \boldsymbol{\Sigma}^{(j)}_{k|k-1}(\mathbf{G}^{-}_{j,k})^{\top}\mathbf{W}^{-1}_{j,k}.\end{aligned} \tag{99}$$

The filtering update is

$$\begin{aligned}\boldsymbol{\eta}^{(j)}_{k|k} &= \boldsymbol{\eta}^{(j)}_{k|k-1} + \mathbf{K}_{j,k}\left(\mathbf{y}^{(j)}_{k} - \mathbf{h}(\boldsymbol{\eta}^{(j)}_{k|k-1})\right),\\ \boldsymbol{\Sigma}^{(j)}_{k|k} &= \boldsymbol{\Sigma}^{(j)}_{k|k-1} - \mathbf{K}_{j,k}\mathbf{W}_{j,k}\mathbf{K}^{\top}_{j,k}.\end{aligned} \tag{100}$$

For numerical implementation, the covariance can equivalently be evaluated in Joseph form:

$$\begin{aligned}\boldsymbol{\Sigma}^{(j)}_{k|k} = (\mathbf{I}_N - \mathbf{K}_{j,k}\mathbf{G}^{-}_{j,k})\boldsymbol{\Sigma}^{(j)}_{k|k-1}(\mathbf{I}_N - \mathbf{K}_{j,k}\mathbf{G}^{-}_{j,k})^{\top}\\ + \mathbf{K}_{j,k}\boldsymbol{\Sigma}^{(r)}_{v}\mathbf{K}^{\top}_{j,k}.\end{aligned}$$

The observation update uses a first-order linearization. The quadratic moment formulas in Lemma 1 are used subsequently to evaluate the observation residual second moments in the conditional maximization step.

RTS backward pass. We initialize the RTS smoother [46], [58] with the terminal filtering moments. For $k = N_t - 1, \ldots, 0$, compute

$$\begin{aligned}\mathbf{A}^{\mathrm{s}}_{j,k} &= \boldsymbol{\Sigma}^{(j)}_{k|k}(\mathbf{F}^{(r)}_{j,k})^{\top}(\boldsymbol{\Sigma}^{(j)}_{k+1|k})^{-1},\\ \boldsymbol{\eta}^{(j)}_{k|N_t} &= \boldsymbol{\eta}^{(j)}_{k|k} + \mathbf{A}^{\mathrm{s}}_{j,k}\left(\boldsymbol{\eta}^{(j)}_{k+1|N_t} - \boldsymbol{\eta}^{(j)}_{k+1|k}\right),\\ \boldsymbol{\Sigma}^{(j)}_{k|N_t} &= \boldsymbol{\Sigma}^{(j)}_{k|k} + \mathbf{A}^{\mathrm{s}}_{j,k}\left(\boldsymbol{\Sigma}^{(j)}_{k+1|N_t} - \boldsymbol{\Sigma}^{(j)}_{k+1|k}\right)(\mathbf{A}^{\mathrm{s}}_{j,k})^{\top},\\ \boldsymbol{\Sigma}^{(j)}_{k+1,k|N_t} &= \boldsymbol{\Sigma}^{(j)}_{k+1|N_t}(\mathbf{A}^{\mathrm{s}}_{j,k})^{\top}.\end{aligned} \tag{101}$$

Here,

$$\boldsymbol{\Sigma}^{(j)}_{k+1,k|N_t} := \mathrm{Cov}_{q^{(r,j)}_{\mathrm{s}}}(\mathbf{z}^{(j)}_{k+1}, \mathbf{z}^{(j)}_{k}).$$

To see the origin of these recursions, the Gaussian backward conditional has mean

$$\boldsymbol{\eta}^{(j)}_{k|k} + \mathbf{A}^{\mathrm{s}}_{j,k}(\mathbf{z}^{(j)}_{k+1} - \boldsymbol{\eta}^{(j)}_{k+1|k})$$

and covariance

$$\mathbf{D}^{\mathrm{s}}_{j,k} := \boldsymbol{\Sigma}^{(j)}_{k|k} - \mathbf{A}^{\mathrm{s}}_{j,k}\boldsymbol{\Sigma}^{(j)}_{k+1|k}(\mathbf{A}^{\mathrm{s}}_{j,k})^{\top}.$$

Combining these backward conditionals with the terminal filtering density defines a joint Gaussian distribution over the latent trajectory. The laws of total expectation and covariance give the smoothed marginal moments. The lag-one covariance follows from the affine dependence of the backward conditional mean on $\mathbf{z}^{(j)}_{k+1}$.

Thus, the E-step returns a coherent joint Gaussian approximation, including the temporal cross-covariances needed for learning. It does not compute the exact posterior of the nonlinear observation model. All matrix inverses in the recursions are implemented through linear solves.

D.4 Proof of Lemma 1

The reconstruction contains the bilinear products $\mathrm{vec}^{-1}(\mathbf{Z}_0)\mathbf{P}_0$ and $\mathrm{vec}^{-1}(\mathbf{Z}_0)\mathbf{Y}_0$, together with linear selections of $\mathbf{Z}_0$ and $\boldsymbol{\mu}_1$. Hence every component h_ℓ is at most quadratic, with constant symmetric Hessian $\mathbf{H}_\ell$.

Let $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\eta}, \boldsymbol{\Sigma})$ and $\boldsymbol{\delta} := \mathbf{z} - \boldsymbol{\eta}$. Writing $\mathbf{g}_\ell := \nabla h_\ell(\boldsymbol{\eta})$, the exact Taylor expansion is

$$h_\ell(\mathbf{z}) = h_\ell(\boldsymbol{\eta}) + \mathbf{g}_\ell^{\top}\boldsymbol{\delta} + \frac{1}{2}\boldsymbol{\delta}^{\top}\mathbf{H}_\ell\boldsymbol{\delta}.$$

Since $\mathbb{E}[\boldsymbol{\delta}] = \mathbf{0}$ and $\mathbb{E}[\boldsymbol{\delta}\boldsymbol{\delta}^{\top}] = \boldsymbol{\Sigma}$,

$$\mathbb{E}[h_\ell(\mathbf{z})] = h_\ell(\boldsymbol{\eta}) + \frac{1}{2}\,\mathrm{tr}(\mathbf{H}_\ell\boldsymbol{\Sigma}).$$

The centered reconstruction component is

$$h_\ell(\mathbf{z}) - \mathbb{E}[h_\ell(\mathbf{z})] = \mathbf{g}_\ell^{\top}\boldsymbol{\delta} + \frac{1}{2}\left(\boldsymbol{\delta}^{\top}\mathbf{H}_\ell\boldsymbol{\delta} - \mathrm{tr}(\mathbf{H}_\ell\boldsymbol{\Sigma})\right).$$

The mixed linear–quadratic contributions to the covariance vanish because all centered Gaussian third moments are zero. For the quadratic contributions, the Gaussian fourth-moment identity gives

$$\begin{aligned}&\mathbb{E}\Big[(\boldsymbol{\delta}^\top \mathbf{H}_\ell \boldsymbol{\delta})(\boldsymbol{\delta}^\top \mathbf{H}_s \boldsymbol{\delta})\Big] \\ &= \operatorname{tr}(\mathbf{H}_\ell \boldsymbol{\Sigma}) \operatorname{tr}(\mathbf{H}_s \boldsymbol{\Sigma}) + 2 \operatorname{tr}(\mathbf{H}_\ell \boldsymbol{\Sigma} \mathbf{H}_s \boldsymbol{\Sigma}).\end{aligned}$$

Therefore,

$$\operatorname{Cov}(h_\ell(\mathbf{z}), h_s(\mathbf{z})) = \mathbf{g}_\ell^\top \boldsymbol{\Sigma} \mathbf{g}_s + \frac{1}{2} \operatorname{tr}(\mathbf{H}_\ell \boldsymbol{\Sigma} \mathbf{H}_s \boldsymbol{\Sigma}).$$

Finally, for a fixed observation $\mathbf{y}$, let $\bar{\mathbf{h}} := \mathbb{E}[\mathbf{h}(\mathbf{z})]$ and $\boldsymbol{\Omega}_h := \operatorname{Cov}(\mathbf{h}(\mathbf{z}))$. The residual $\mathbf{e} := \mathbf{y} - \mathbf{h}(\mathbf{z})$ has mean $\mathbf{y} - \bar{\mathbf{h}}$ and covariance $\boldsymbol{\Omega}_h$. Consequently,

$$\mathbb{E}[\mathbf{e}\mathbf{e}^\top] = (\mathbf{y} - \bar{\mathbf{h}})(\mathbf{y} - \bar{\mathbf{h}})^\top + \boldsymbol{\Omega}_h,$$

which proves the residual-moment identity.

The identities are exact for any Gaussian argument distribution. In the conditional covariance update, that distribution is the smoothed marginal produced by EKF–RTS. Thus, the posterior approximation does not require an additional linearization of the reconstruction when evaluating its expected residual second moment. The Hessian contractions involve only the coordinates read by $\mathbf{h}$ and can be evaluated using the corresponding covariance submatrix.

D.5 Expected Residual Second Moments

Throughout this subsection, the smoothing moments obtained at iteration r remain fixed. Candidate physical parameters enter only through $\mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}})$ and $\mathbf{b}_{j,k}(\boldsymbol{\theta}_{\text{phys}})$.

Transition residual. Define its mean as

$$\begin{aligned}\mathbf{d}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) := \boldsymbol{\eta}^{(j)}_{k+1|N_t} - \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) \boldsymbol{\eta}^{(j)}_{k|N_t} \\ - \mathbf{b}_{j,k}(\boldsymbol{\theta}_{\text{phys}}).\end{aligned}$$

Expanding the covariance of the affine residual in (21) gives

$$\begin{aligned}\mathbf{S}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) = \mathbf{d}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) \mathbf{d}_{j,k}(\boldsymbol{\theta}_{\text{phys}})^\top + \boldsymbol{\Sigma}^{(j)}_{k+1|N_t} \\ + \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) \boldsymbol{\Sigma}^{(j)}_{k|N_t} \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}})^\top \\ - \boldsymbol{\Sigma}^{(j)}_{k+1,k|N_t} \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}})^\top \\ - \mathbf{F}_{j,k}(\boldsymbol{\theta}_{\text{phys}}) (\boldsymbol{\Sigma}^{(j)}_{k+1,k|N_t})^\top.\end{aligned} \tag{102}$$

The two cross terms arise from the temporal correlation between consecutive latent states. Replacing this expression by the outer product of the mean residual alone would discard both marginal uncertainty and temporal cross-covariance.

Observation residual. Apply Lemma 1 to the smoothed marginal $\mathcal{N}(\boldsymbol{\eta}^{(j)}_{k|N_t}, \boldsymbol{\Sigma}^{(j)}_{k|N_t})$. Denote the resulting reconstruction mean and covariance by $\bar{\mathbf{h}}^{\text{s}}_{j,k}$ and $\boldsymbol{\Omega}^{\text{s}}_{j,k}$, respectively. Then

$$\mathbf{V}_{j,k} = (\mathbf{y}^{(j)}_k - \bar{\mathbf{h}}^{\text{s}}_{j,k})(\mathbf{y}^{(j)}_k - \bar{\mathbf{h}}^{\text{s}}_{j,k})^\top + \boldsymbol{\Omega}^{\text{s}}_{j,k}. \tag{103}$$

This expression includes uncertainty in the reconstructed observation as well as its mean prediction error. It supplies the sufficient residual statistic for updating $\boldsymbol{\Sigma}_v$.

Initial residual. Since the initialization mean $\boldsymbol{\Phi}(\mathbf{y}^{(j)}_0)$ is fixed,

$$\begin{aligned}\mathbf{S}^0_j = \boldsymbol{\Sigma}^{(j)}_{0|N_t} \\ + (\boldsymbol{\eta}^{(j)}_{0|N_t} - \boldsymbol{\Phi}(\mathbf{y}^{(j)}_0))(\boldsymbol{\eta}^{(j)}_{0|N_t} - \boldsymbol{\Phi}(\mathbf{y}^{(j)}_0))^\top.\end{aligned} \tag{104}$$

All three residual second moments are evaluated exactly under the fixed Gaussian smoothing approximation. In particular, $\mathbf{V}_{j,k}$ and $\mathbf{S}^0_j$ do not depend on candidate parameters during the conditional updates.

D.6 Feasible Physical Parameter Updates

We enforce the physical constraints $m > 0$, $\mathbf{J} \succ 0$, $\mathbf{D}_v \succeq 0$, and $\mathbf{D}_\omega \succeq 0$ through

$$m = \exp(\alpha), \qquad \mathbf{J} = \mathbf{L}_J \mathbf{L}_J^\top, \qquad \mathbf{D}_v = \mathbf{L}_v \mathbf{L}_v^\top, \qquad \mathbf{D}_\omega = \mathbf{L}_\omega \mathbf{L}_\omega^\top.$$

Here, $\alpha \in \mathbb{R}$ is a free scalar, the factors are lower triangular, and the diagonal entries of $\mathbf{L}_J$ are exponentials of free parameters. Let $\boldsymbol{\xi} \in \mathbb{R}^{19}$ collect the resulting free parameters and write $\boldsymbol{\theta}_{\text{phys}} = \mathcal{T}(\boldsymbol{\xi})$. This parametrization preserves feasibility at every iterate without projection.

With the covariance matrices fixed at their current values, define

$$f_r(\boldsymbol{\xi}) := \mathcal{L}_r(\mathcal{T}(\boldsymbol{\xi})),$$

where $\mathcal{L}_r$ is given in (25). We decrease f_r using a modified Newton method with eigenvalue flooring and Armijo backtracking [51]. The objective is evaluated using precomputed E-step sufficient statistics, which remain fixed throughout the physical update.

Modified Newton direction. At each inner iteration, compute $\mathbf{g} := \nabla f_r(\boldsymbol{\xi})$ by reverse-mode automatic differentiation and $\mathbf{H} := \nabla^2 f_r(\boldsymbol{\xi})$ by forward-over-reverse automatic differentiation. Both derivatives account for the analytical construction and discretization of the transition matrices; no differentiation is performed through the E-step.

Symmetrize the computed Hessian and write

$$\frac{1}{2}(\mathbf{H} + \mathbf{H}^\top) = \mathbf{U}\,\mathrm{diag}(\lambda_i)\mathbf{U}^\top.$$

For a prescribed floor $\delta_H > 0$, define

$$\widetilde{\mathbf{H}} := \mathbf{U}\,\mathrm{diag}\left(\max\{\lambda_i, \delta_H\}\right)\mathbf{U}^\top, \qquad \mathbf{p} := -\widetilde{\mathbf{H}}^{-1}\mathbf{g}. \tag{105}$$

The direction is computed by solving $\widetilde{\mathbf{H}}\mathbf{p} = -\mathbf{g}$. Since $\widetilde{\mathbf{H}} \succ 0$, for $\mathbf{g} \neq \mathbf{0}$,

$$\mathbf{g}^\top\mathbf{p} = -\mathbf{g}^\top\widetilde{\mathbf{H}}^{-1}\mathbf{g} < 0.$$

Thus, eigenvalue flooring supplies a descent direction even when the physical objective has an indefinite Hessian. If the computed direction fails the descent check numerically, we use the fallback $\mathbf{p} = -\mathbf{P}\mathbf{g}$, where $\mathbf{P} \succ 0$ is the diagonal preconditioner determined by the parameter scales.

Step acceptance. Backtracking considers $\gamma = \gamma_0\beta^a$, with $\gamma_0 > 0$, $\beta = 0.5$, and $a \in \mathbb{N}_0$. A trial update is accepted only if its objective value is finite and satisfies

$$f_r(\boldsymbol{\xi} + \gamma\mathbf{p}) \leq f_r(\boldsymbol{\xi}) + c_{\text{A}}\gamma\mathbf{g}^\top\mathbf{p}, \qquad c_{\text{A}} = 10^{-4}. \tag{106}$$

On acceptance, set $\boldsymbol{\xi} \leftarrow \boldsymbol{\xi} + \gamma\mathbf{p}$. The initial trial step for the next inner iteration is increased by a factor $1/\beta$, subject to the prescribed relative-step cap of 1.0. If the gradient vanishes or no acceptable step is found within the backtracking budget, the current parameters are retained. We allow at most twelve accepted Newton steps per physical update. Consequently, every accepted step decreases f_r, and the resulting physical parameters satisfy $\mathcal{L}_r(\boldsymbol{\theta}^{(r+1)}_{\text{phys}}) \leq \mathcal{L}_r(\boldsymbol{\theta}^{(r)}_{\text{phys}})$.

The initial physical parameters are feasible, and the initial covariance matrices satisfy their prescribed positive eigenvalue floors. The drag factors are initialized away from the all-zero matrix because the derivative of $\mathbf{L}\mathbf{L}^\top$ vanishes at $\mathbf{L} = \mathbf{0}$. The parametrization and line search ensure feasibility and objective decrease, respectively; they do not imply convexity or global optimality.

The covariance matrices are excluded from this Newton iteration. After the physical update, $\boldsymbol{\Sigma}_w$, $\boldsymbol{\Sigma}_v$, and $\boldsymbol{\Sigma}_0$ are updated by the closed-form conditional maximizers in (26).

D.7 Proof of Lemma 2

Proof of Lemma 2. Let $\mathbf{S} \succeq 0$ and $\epsilon > 0$. Every feasible covariance $\boldsymbol{\Sigma} \succeq \epsilon\mathbf{I}_d$ is positive definite. We may therefore introduce its precision matrix $\mathbf{X} := \boldsymbol{\Sigma}^{-1}$. Since inversion replaces each positive eigenvalue by its reciprocal, the covariance constraint is equivalent to

$$\boldsymbol{\Sigma} \succeq \epsilon\mathbf{I}_d \quad \Longleftrightarrow \quad 0 \prec \mathbf{X} \preceq \epsilon^{-1}\mathbf{I}_d.$$

Moreover,

$$\log\det\boldsymbol{\Sigma} = \log\det(\mathbf{X}^{-1}) = -\log\det\mathbf{X}, \qquad \mathrm{tr}(\boldsymbol{\Sigma}^{-1}\mathbf{S}) = \mathrm{tr}(\mathbf{X}\mathbf{S}).$$

Thus, the negative sign arises from the determinant of the inverse. The substitution preserves the objective value, and the original minimization becomes

$$\min_{0 \prec \mathbf{X} \preceq \epsilon^{-1}\mathbf{I}_d} \left\{-\log\det\mathbf{X} + \mathrm{tr}(\mathbf{X}\mathbf{S})\right\}.$$

Let $\mathbf{S} = \mathbf{U}\,\mathrm{diag}(s_1, \ldots, s_d)\mathbf{U}^\top$ be an eigendecomposition, where $\mathbf{U}$ is orthogonal and $s_i \geq 0$. Express the precision matrix in this eigenbasis as $\mathbf{Y} := \mathbf{U}^\top\mathbf{X}\mathbf{U}$. Orthogonal similarity preserves the

determinant and the eigenvalue constraints, while

$$\operatorname{tr}(\mathbf{XS}) = \sum_{i=1}^{d} s_i Y_{ii}.$$

Define $\mathbf{D} := \operatorname{diag}(Y_{11}, \ldots, Y_{dd})$. Since $0 < Y_{ii} \leq \epsilon^{-1}$, the matrix $\mathbf{D}$ is feasible and has the same trace contribution as $\mathbf{Y}$. Hadamard's inequality gives

$$\det \mathbf{Y} \leq \prod_{i=1}^{d} Y_{ii} = \det \mathbf{D},$$

and hence $-\log\det \mathbf{D} \leq -\log\det \mathbf{Y}$. Replacing $\mathbf{Y}$ by $\mathbf{D}$ therefore cannot increase the objective. It suffices to optimize over precision matrices that are diagonal in the eigenbasis of $\mathbf{S}$.

Writing the corresponding covariance as $\mathbf{\Sigma} = \mathbf{U}\operatorname{diag}(\lambda_1, \ldots, \lambda_d)\mathbf{U}^\top$, with $\lambda_i \geq \epsilon$, reduces the problem to the independent scalar minimizations

$$\min_{\lambda_i \geq \epsilon} \left\{ \log \lambda_i + \frac{s_i}{\lambda_i} \right\}, \qquad i = 1, \ldots, d.$$

For each scalar objective,

$$\frac{d}{d\lambda_i}\left(\log \lambda_i + \frac{s_i}{\lambda_i}\right) = \frac{\lambda_i - s_i}{\lambda_i^2}.$$

If $s_i \leq \epsilon$, the objective is minimized at $\lambda_i = \epsilon$. If $s_i > \epsilon$, it decreases up to $\lambda_i = s_i$ and increases thereafter. Consequently,

$$\lambda_i^\star = \max\{s_i, \epsilon\},$$

which yields

$$\mathbf{\Sigma}^\star = \mathbf{U}\operatorname{diag}\big(\max\{s_1, \epsilon\}, \ldots, \max\{s_d, \epsilon\}\big)\mathbf{U}^\top = \mathcal{P}_\epsilon(\mathbf{S}).$$

Finally, the precision objective $-\log\det \mathbf{X} + \operatorname{tr}(\mathbf{XS})$ is strictly convex on its convex feasible set. Its minimizer is therefore unique, and the bijective relation $\mathbf{\Sigma} = \mathbf{X}^{-1}$ establishes uniqueness of the covariance minimizer. □

D.8 CM-step 2 Derivation

After CM-step 1, fix the physical parameters at $\boldsymbol{\theta}_{\text{phys}}^{(r+1)}$ and retain the smoothing distributions $\{q_{\text{s}}^{(r,j)}\}_{j=1}^{M}$ computed in the E-step. The residual second moments are therefore fixed with respect to the candidate covariance matrices. Define their averages as

$$\bar{\mathbf{S}}_w := \tfrac{1}{MN_t}\sum_{j=1}^{M}\sum_{k=0}^{N_t-1} \mathbf{S}_{j,k}(\boldsymbol{\theta}_{\text{phys}}^{(r+1)}),\ \bar{\mathbf{S}}_v := \frac{1}{MN_t}\sum_{j=1}^{M}\sum_{k=1}^{N_t} \mathbf{V}_{j,k},\ \bar{\mathbf{S}}_0 := \frac{1}{M}\sum_{j=1}^{M} \mathbf{S}_j^0.$$

Each trajectory contributes N_t transition residuals, N_t observation residuals, and one initial residual, which explains the respective normalization factors MN_t, MN_t, and M. Only $\bar{\mathbf{S}}_w$ depends on the newly updated physical parameters; all three averages use the same fixed E-step distributions.

The covariance update maximizes $\mathcal{Q}_r$ with the physical parameters held fixed. Equivalently, it minimizes $-\mathcal{Q}_r$, since negating an objective reverses the ordering of its values. Using linearity of the trace, (95) becomes

$$\begin{aligned} -\mathcal{Q}_r(\boldsymbol{\theta}) = c &+ \frac{MN_t}{2}\left[\log\det \mathbf{\Sigma}_w + \operatorname{tr}\left(\mathbf{\Sigma}_w^{-1}\bar{\mathbf{S}}_w\right)\right] \\ &+ \frac{MN_t}{2}\left[\log\det \mathbf{\Sigma}_v + \operatorname{tr}\left(\mathbf{\Sigma}_v^{-1}\bar{\mathbf{S}}_v\right)\right] \\ &+ \frac{M}{2}\left[\log\det \mathbf{\Sigma}_0 + \operatorname{tr}\left(\mathbf{\Sigma}_0^{-1}\bar{\mathbf{S}}_0\right)\right], \end{aligned}$$

where c is independent of the covariance matrices.

Each covariance appears in a separate term, and the constraints

$$\mathbf{\Sigma}_w \succeq \epsilon_w \mathbf{I}_N, \qquad \mathbf{\Sigma}_v \succeq \epsilon_v \mathbf{I}_{18}, \qquad \mathbf{\Sigma}_0 \succeq \epsilon_0 \mathbf{I}_N$$

also separate across the three blocks. Consequently, the covariance update reduces to three independent minimization problems. The additive constant c can be discarded, and dividing each term by its

strictly positive factor $MN_t/2$ or $M/2$ leaves its minimizer unchanged. Thus,

$$\boldsymbol{\Sigma}_w^{(r+1)} = \operatorname*{argmin}_{\boldsymbol{\Sigma}_w \succeq \epsilon_w \mathbf{I}_N} \left\{ \log\det \boldsymbol{\Sigma}_w + \operatorname{tr}\left(\boldsymbol{\Sigma}_w^{-1} \bar{\mathbf{S}}_w\right) \right\},$$
$$\boldsymbol{\Sigma}_v^{(r+1)} = \operatorname*{argmin}_{\boldsymbol{\Sigma}_v \succeq \epsilon_v \mathbf{I}_{18}} \left\{ \log\det \boldsymbol{\Sigma}_v + \operatorname{tr}\left(\boldsymbol{\Sigma}_v^{-1} \bar{\mathbf{S}}_v\right) \right\},$$
$$\boldsymbol{\Sigma}_0^{(r+1)} = \operatorname*{argmin}_{\boldsymbol{\Sigma}_0 \succeq \epsilon_0 \mathbf{I}_N} \left\{ \log\det \boldsymbol{\Sigma}_0 + \operatorname{tr}\left(\boldsymbol{\Sigma}_0^{-1} \bar{\mathbf{S}}_0\right) \right\}.$$

Each residual second moment is the expectation of an outer product and is therefore positive semidefinite. Their averages $\bar{\mathbf{S}}_w$, $\bar{\mathbf{S}}_v$, and $\bar{\mathbf{S}}_0$ are consequently positive semidefinite as well. Applying Lemma 2 to the three problems gives

$$\boldsymbol{\Sigma}_w^{(r+1)} = \mathcal{P}_{\epsilon_w}(\bar{\mathbf{S}}_w), \qquad \boldsymbol{\Sigma}_v^{(r+1)} = \mathcal{P}_{\epsilon_v}(\bar{\mathbf{S}}_v), \qquad \boldsymbol{\Sigma}_0^{(r+1)} = \mathcal{P}_{\epsilon_0}(\bar{\mathbf{S}}_0).$$

Substituting the definitions of the averages recovers (26). These updates are the unique conditional maximizers over the prescribed covariance feasible sets. The initial mean remains $\boldsymbol{\Phi}(\mathbf{y}_0^{(j)})$ for each trajectory.

If a covariance is constrained to be diagonal, its update instead has diagonal entries $\max\{\bar{S}_{ii}, \epsilon\}$ and zero off-diagonal entries. For a prescribed block-diagonal covariance, the objective separates over the specified blocks, and Lemma 2 applies to each corresponding principal block of the average residual matrix. The covariance structures and positive floors remain fixed throughout training.

D.9 Proof of Conditional Surrogate Improvement

Proof of Theorem 1. Fix the smoothing distributions $\{q_{\mathrm{s}}^{(r,j)}\}_{j=1}^M$ throughout both conditional parameter updates. Suppose that $\boldsymbol{\theta}^{(r)}$ satisfies the physical constraints and covariance floors.

With the covariance matrices fixed at their iteration-r values, (95) gives

$$-\mathcal{Q}_r\left(\boldsymbol{\theta}_{\mathrm{phys}}, \boldsymbol{\Sigma}_w^{(r)}, \boldsymbol{\Sigma}_v^{(r)}, \boldsymbol{\Sigma}_0^{(r)}\right) = C_r + \mathcal{L}_r(\boldsymbol{\theta}_{\mathrm{phys}}),$$

where C_r is independent of the candidate physical parameters. Indeed, the physical parameters enter only through the transition residual moments, whose contribution is precisely $\mathcal{L}_r$. All remaining terms are fixed during CM-step 1.

Define the intermediate parameter vector

$$\boldsymbol{\theta}^{\mathrm{int}} := \left(\boldsymbol{\theta}_{\mathrm{phys}}^{(r+1)}, \boldsymbol{\Sigma}_w^{(r)}, \boldsymbol{\Sigma}_v^{(r)}, \boldsymbol{\Sigma}_0^{(r)}\right).$$

The acceptance condition $\mathcal{L}_r(\boldsymbol{\theta}_{\mathrm{phys}}^{(r+1)}) \le \mathcal{L}_r(\boldsymbol{\theta}_{\mathrm{phys}}^{(r)})$ therefore implies

$$\mathcal{Q}_r(\boldsymbol{\theta}^{\mathrm{int}}) - \mathcal{Q}_r(\boldsymbol{\theta}^{(r)}) = \mathcal{L}_r(\boldsymbol{\theta}_{\mathrm{phys}}^{(r)}) - \mathcal{L}_r(\boldsymbol{\theta}_{\mathrm{phys}}^{(r+1)}) \ge 0.$$

Thus, decreasing the physical objective increases the expected complete-data objective. If no physical update is accepted, the current parameters are retained and this difference is zero. The physical parameterization preserves the constraints, while the unchanged covariance matrices remain feasible.

For CM-step 2, hold the physical parameters fixed at $\boldsymbol{\theta}_{\mathrm{phys}}^{(r+1)}$. As established in § D.8, the covariance-dependent terms separate across $\boldsymbol{\Sigma}_w$, $\boldsymbol{\Sigma}_v$, and $\boldsymbol{\Sigma}_0$. By Lemma 2, each update in (26) minimizes its corresponding contribution to $-\mathcal{Q}_r$ over the prescribed feasible set. Consequently, the three updates jointly maximize $\mathcal{Q}_r$ over the covariance blocks.

The previous covariance matrices are admissible candidates for this optimization because they satisfy the same fixed floors. The updated covariances therefore satisfy

$$\mathcal{Q}_r(\boldsymbol{\theta}^{(r+1)}) \ge \mathcal{Q}_r(\boldsymbol{\theta}^{\mathrm{int}}).$$

Moreover, each spectral update has eigenvalues at least equal to its prescribed positive floor, so covariance feasibility is preserved. Combining the two conditional updates yields

$$\mathcal{Q}_r(\boldsymbol{\theta}^{(r+1)}) \ge \mathcal{Q}_r(\boldsymbol{\theta}^{\mathrm{int}}) \ge \mathcal{Q}_r(\boldsymbol{\theta}^{(r)}).$$

Finally, the fixed-distribution lower bound in (97) is

$$\mathcal{B}_r(\boldsymbol{\theta}) = \mathcal{Q}_r(\boldsymbol{\theta}) + \sum_{j=1}^M \mathcal{H}(q_{\mathrm{s}}^{(r,j)}).$$

Since the smoothing distributions remain fixed, their entropy sum is independent of the candidate parameters. Hence,

$$\mathcal{B}_r(\boldsymbol{\theta}^{(r+1)}) - \mathcal{B}_r(\boldsymbol{\theta}^{(r)}) = \mathcal{Q}_r(\boldsymbol{\theta}^{(r+1)}) - \mathcal{Q}_r(\boldsymbol{\theta}^{(r)}) \ge 0.$$

Algorithm 1: GECM Training of ABS-PWM

Input: Data $\mathcal{D}$; lifting degree K; feasible initialization $\boldsymbol{\theta}^{(0)}$; $\tau_{\mathrm{LL}} = 2 \times 10^{-5}$; patience $p = 5$; $R_{\max} = 200$; $I_{\max} = 12$.
Output: Learned parameters $\boldsymbol{\theta}^{\star}$.

$c_{\mathrm{LL}} \leftarrow 0$
for $r = 0, \ldots, R_{\max} - 1$ **do**
  Construct the transition matrices from $\boldsymbol{\theta}^{(r)}_{\mathrm{phys}}$
  `// E-step`
  for $j = 1, \ldots, M$ **do**
    $\boldsymbol{\eta}^{(j)}_{0|0} \leftarrow \boldsymbol{\Phi}(\mathbf{y}^{(j)}_0)$, $\boldsymbol{\Sigma}^{(j)}_{0|0} \leftarrow \boldsymbol{\Sigma}^{(r)}_0$
    Run EKF–RTS using (98)–(101)
    Compute $\mathbf{V}_{j,k}$ and $\mathbf{S}^0_j$ using (103) and (104)
  $\ell_r \leftarrow \widehat{\mathfrak{L}}_{\mathrm{EKF}}(\boldsymbol{\theta}^{(r)})$ using the EKF predictive moments
  Hold the smoothing distributions fixed throughout both CM-steps
  `// CM-step 1`
  Update $\boldsymbol{\theta}^{(r+1)}_{\mathrm{phys}}$ using at most $I_{\max}$ feasible modified Newton iterations for (25), with Armijo acceptance (106)
  Retain the current physical iterate if no acceptable step is found
  `// CM-step 2`
  Evaluate $\mathbf{S}_{j,k}(\boldsymbol{\theta}^{(r+1)}_{\mathrm{phys}})$ using (102)
  Update $\boldsymbol{\Sigma}^{(r+1)}_w$, $\boldsymbol{\Sigma}^{(r+1)}_v$, and $\boldsymbol{\Sigma}^{(r+1)}_0$ using (26), and assemble $\boldsymbol{\theta}^{(r+1)}$
  `// Predictive log-likelihood stabilization`
  if $r > 0$ **then**
    $d_{\mathrm{LL}} \leftarrow \dfrac{|\ell_r - \ell_{r-1}|}{\max\{|\ell_{r-1}|, \epsilon_{\mathrm{num}}\}}$
    if $d_{\mathrm{LL}} < \tau_{\mathrm{LL}}$ **then**
      $c_{\mathrm{LL}} \leftarrow c_{\mathrm{LL}} + 1$
    else
      $c_{\mathrm{LL}} \leftarrow 0$
    if $c_{\mathrm{LL}} \geq p$ **then**
      return $\boldsymbol{\theta}^{(r+1)}$
return $\boldsymbol{\theta}^{(R_{\max})}$

This establishes feasibility preservation and nondecrease of both the surrogate objective and its associated fixed-distribution lower bound. □

Remark 2 (EKF predictive log-likelihood)**.** *The conditional observed-data log-likelihood admits the exact factorization*

$$\mathfrak{L}(\boldsymbol{\theta}) = \sum_{j=1}^{M} \sum_{k=1}^{N_t} \log p_{\boldsymbol{\theta}}\left(\mathbf{y}^{(j)}_k \mid \mathbf{y}^{(j)}_{0:k-1}, \mathbf{u}^{(j)}_{0:k-1}\right).$$

Writing the conditioning history as $\mathcal{H}_{j,k-1}$, each predictive density is

$$p_{\boldsymbol{\theta}}(\mathbf{y}^{(j)}_k \mid \mathcal{H}_{j,k-1}) = \int \mathcal{N}\left(\mathbf{y}^{(j)}_k; h(\mathbf{z}), \boldsymbol{\Sigma}_v\right) p_{\boldsymbol{\theta}}(\mathbf{z} \mid \mathcal{H}_{j,k-1})\, \mathrm{d}\mathbf{z}.$$

The EKF equations (98)–(100) approximate the predictive latent distribution by a Gaussian and linearize h at $\boldsymbol{\eta}^{(j)}_{k|k-1}$, giving

$$p_{\boldsymbol{\theta}}(\mathbf{y}^{(j)}_k \mid \mathcal{H}_{j,k-1}) \approx \mathcal{N}\left(\mathbf{y}^{(j)}_k; h(\boldsymbol{\eta}^{(j)}_{k|k-1}), \mathbf{W}_{j,k}\right).$$

Consequently, the EKF predictive log-likelihood estimate is

$$\widehat{\mathfrak{L}}_{\mathrm{EKF}}(\boldsymbol{\theta}) := -\frac{1}{2} \sum_{j=1}^{M} \sum_{k=1}^{N_t} \left[18 \log(2\pi) + \log\det \mathbf{W}_{j,k} + \mathbf{e}^{\top}_{j,k} \mathbf{W}^{-1}_{j,k} \mathbf{e}_{j,k}\right],$$

where $\mathbf{e}_{j,k}$ and $\mathbf{W}_{j,k}$ are the EKF innovation and its covariance, computed under $\boldsymbol{\theta}$ before incorporating $\mathbf{y}_k^{(j)}$. This score approximates $\mathfrak{L}(\boldsymbol{\theta})$ and is generally neither a lower nor an upper bound. Its stabilization is used as the stopping criterion in Algorithm 1.

E SUPPLEMENTARY MATERIAL FOR § 5

This appendix develops the prediction model, covariance weighting, and preliminary theoretical properties of the controller in § 5. The ESKF supplies a physical-state estimate, the MPC computes a nominal input plan, and a separate observer estimates the wrench correction.

E.1 STATE ESTIMATION AND MEASUREMENT CONVENTIONS

The estimate $\hat{\mathbf{x}}_{k|k}$ is obtained from measurements through time k using a standard ESKF with right attitude errors [52]. Its nominal rigid-body predictor uses nominal mass and inertia, gravity, and the previously applied wrench, with no drag or mixer correction. The filter runs at Δt and uses the sensor model $\mathbf{y}_k = \mathbf{x}_k \boxplus \boldsymbol{\nu}_k$, calibrated tangent covariance $\boldsymbol{\Sigma}_\nu \in \mathbb{R}^{12\times 12}$, and prescribed per-sample process covariance $\mathbf{Q}_{\text{est}} \succeq 0$. It is initialized from the first measurement with error covariance $\boldsymbol{\Sigma}_\nu$, without counting that measurement twice. Joseph-form covariance correction and the attitude-error reset follow [52, Sec. 6]; an identity reset is a small-correction approximation.

The ESKF maintains a physical attitude, so $\boldsymbol{\Phi}(\hat{\mathbf{x}}_{k|k})$ is a consistent physical lift for MPC initialization. The separate wrench observer uses $\mathbf{z}_k^{\text{m}} = \boldsymbol{\Phi}(\mathbf{y}_k)$ and the learned ABS-PWM covariances $\boldsymbol{\Sigma}_w$ and $\boldsymbol{\Sigma}_v$. These are not the process and measurement covariances of the nominal-model ESKF. In particular, the learned ambient $\boldsymbol{\Sigma}_v \in \mathbb{R}^{18\times 18}$ and tangent sensor covariance $\boldsymbol{\Sigma}_\nu \in \mathbb{R}^{12\times 12}$ are distinct quantities.

E.2 GUIDE TRAJECTORIES AND THE TWO PREDICTION INTERVALS

The guide is a fixed trajectory used to evaluate the state-dependent input map; the reference specifies the desired motion in the objective. Reference states can initialize the guide. Subsequently, the implemented warm start shifts the previous optimal predictions by one shooting interval and reconstructs and re-lifts their states after projecting the guide attitude onto $\mathrm{SO}(3)$. This shift is a guide-generation heuristic: when $\delta > \Delta t$, it is not a time-aligned shift by the elapsed sampling period. It supplies fixed coefficients to the new QP and makes no claim of recursive feasibility. The initial condition instead uses $\boldsymbol{\Phi}(\hat{\mathbf{x}}_{k|k})$. The raw measurement-to-lift map used by the wrench observer has no additional attitude projection.

Holding $\mathbf{u}_{l|k}$ and $\mathbf{H}(\mathbf{z}^{\text{g}}_{l|k})$ constant over a shooting interval gives (28) by integrating $\dot{\mathbf{z}} = \mathbf{A}_{\text{aug}}\mathbf{z} + \mathbf{H}(\mathbf{z}^{\text{g}}_{l|k})\mathbf{u}_{l|k}$. The matrices can be computed together from

$$\exp\left(\delta \begin{bmatrix} \mathbf{A}_{\text{aug}} & \mathbf{I}_N \\ \mathbf{0} & \mathbf{0} \end{bmatrix}\right) = \begin{bmatrix} \mathbf{A}_d & \boldsymbol{\Gamma}_\delta \\ \mathbf{0} & \mathbf{I}_N \end{bmatrix}. \tag{107}$$

This is exact for the frozen-guide differential equation, not for the original bilinear trajectory with a varying state-dependent input map.

The observer instead uses the identified transition over Δt:

$$\begin{aligned} f_\Delta(\mathbf{z},\mathbf{u}) &:= \mathbf{A}_\Delta \mathbf{z} + \bar{\mathbf{B}}_\Delta \mathbf{u} + \sum_{i=1}^{4} u_i \mathbf{B}_{\Delta i}\mathbf{z}, \\ \mathbf{G}_\Delta(\mathbf{z}) &:= \bar{\mathbf{B}}_\Delta + [\mathbf{B}_{\Delta 1}\mathbf{z} \quad \cdots \quad \mathbf{B}_{\Delta 4}\mathbf{z}], \qquad \mathbf{F}(\mathbf{u}) := \mathbf{A}_\Delta + \sum_{i=1}^{4} u_i \mathbf{B}_{\Delta i}. \end{aligned} \tag{108}$$

The shooting maps use the continuous augmented drift, while the identified transition uses the manuscript's discrete approximation. In general, $\mathbf{A}_d \neq \mathbf{A}_\Delta^{n_s}$. The observer residual is evaluated with f_Δ because $\boldsymbol{\Sigma}_w$ was identified for that one-sample model.

E.3 CONDENSING, CONVEXITY, AND TERMINAL WEIGHT

Let $\mathbf{E}_l \in \mathbb{R}^{4N_h\times 4}$ insert input block l, so that $\mathbf{u}_{l|k} = \mathbf{E}_l^\top \mathbf{U}_k$. The nominal state is affine in the input sequence:

$$\bar{\mathbf{z}}_{l|k} = \mathbf{v}_{l|k} + \mathbf{S}^z_{l|k}\mathbf{U}_k, \qquad \mathbf{v}_{0|k} = \boldsymbol{\Phi}(\hat{\mathbf{x}}_{k|k}), \quad \mathbf{S}^z_{0|k} = \mathbf{0}, \tag{109}$$

with

$$\mathbf{v}_{l+1|k} = \mathbf{A}_d \mathbf{v}_{l|k}, \qquad \mathbf{S}^z_{l+1|k} = \mathbf{A}_d \mathbf{S}^z_{l|k} + \mathbf{B}_{d,l|k}\mathbf{E}_l^\top. \tag{110}$$

Stack the maps for $l \in [1 : N_h]$, define $\mathcal{Q} := \mathrm{blkdiag}(\mathbf{Q},\ldots,\mathbf{Q},\mathbf{Q}_f)$ and $\mathcal{R} := \mathbf{I}_{N_h} \otimes \mathbf{R}_u$, and use $\mathcal{Q} = \mathbf{Q}_f$ when $N_h = 1$. The Hessian is

$$\mathcal{H} = 2\big((\mathbf{S}^z)^\top \mathcal{Q}\mathbf{S}^z + \mathcal{R}\big) \succ 0. \tag{111}$$

Proposition 2 (Unique nominal input plan)**.** *For fixed guides and guard bounds, positive-semidefinite tracking weights, and* $\mathbf{R}_u \succ 0$*, (29) has a unique minimizing input sequence whenever its feasible set is nonempty.*

Proof. The condensed state is affine in $\mathbf{U}_k$, so all constraints are affine. Equation (111) makes the objective strictly convex and coercive. The feasible set is closed; if nonempty it admits a minimizer, which is unique by strict convexity. □

The terminal weight may be constructed using the virtual-input matrix $\mathbf{E}_U$ of the lifted realization. Let $\mathbf{B}_v := \mathbf{\Gamma}_\delta \mathbf{E}_U$ and $\rho_{\mathrm{term}} > 0$. When a stabilizing positive-semidefinite solution exists, use

$$\begin{aligned}\mathbf{Q}_f = \mathbf{Q} &+ \mathbf{A}_d^\top \mathbf{Q}_f \mathbf{A}_d \\ &- \mathbf{A}_d^\top \mathbf{Q}_f \mathbf{B}_v (\rho_{\mathrm{term}} \mathbf{I} + \mathbf{B}_v^\top \mathbf{Q}_f \mathbf{B}_v)^{-1} \mathbf{B}_v^\top \mathbf{Q}_f \mathbf{A}_d.\end{aligned} \tag{112}$$

This construction requires the usual stabilizability and detectability conditions for the auxiliary pair and cost. The virtual inputs are used only to select a terminal penalty; they are not MPC decision variables. A prescribed positive-semidefinite terminal weight, including $\mathbf{Q}$, can be used if this construction is unavailable. The auxiliary Riccati solution does not certify stability of the physical four-input closed loop.

E.4 Nominal guards and precomputed uncertainty margins

The implemented guard selector extracts $(\mathbf{P}_0, \mathbf{Y}_0, \boldsymbol{\mu}_1)$. At guide attitude $\mathbf{R}^{\mathrm{g}}_{l|k}$, the untightened bounds are symmetric, with positive vector

$$\mathbf{b}^0_{l|k} := \mathrm{col}\Big(|(\mathbf{R}^{\mathrm{g}}_{l|k})^\top| \mathbf{s}_{\max},\ |(\mathbf{R}^{\mathrm{g}}_{l|k})^\top| \mathbf{v}_{\max},\ \boldsymbol{\omega}_{\max} \Big). \tag{113}$$

Matrix absolute values are entrywise. The position and velocity bounds are axis-aligned outer boxes of the rotated physical boxes. Feasibility of these outer boxes does not imply feasibility of the original physical limits, even at the guide attitude. The body-rate selection is exact on physical lifts because $\boldsymbol{\mu}_1 = \boldsymbol{\omega}$.

To construct the nominal uncertainty margins, form

$$\mathbf{W}_\delta := \sum_{r=0}^{n_s-1} \mathbf{A}^r_\Delta \mathbf{\Sigma}_w (\mathbf{A}^r_\Delta)^\top, \qquad \mathbf{\Pi}_0 := \mathbf{\Sigma}_0, \quad \mathbf{\Pi}_{l+1} := \mathbf{A}_d \mathbf{\Pi}_l \mathbf{A}_d^\top + \mathbf{W}_\delta. \tag{114}$$

Let $\sigma_{l,a} := \sqrt{[\mathbf{C}\mathbf{\Pi}_l \mathbf{C}^\top]_{aa}}$. For a fixed margin multiplier $\kappa \geq 0$, set

$$\underline{\mathbf{b}}_{l|k} := -\mathbf{b}^0_{l|k} + \kappa \boldsymbol{\sigma}_l, \qquad \overline{\mathbf{b}}_{l|k} := \mathbf{b}^0_{l|k} - \kappa \boldsymbol{\sigma}_l, \quad l \in [1 : N_h - 1]. \tag{115}$$

The covariance proxies and marginal standard deviations can be precomputed for a fixed model and horizon. This recursion uses only drift matrices, mixes the identification-step accumulation with the shooting drift, and restarts from $\mathbf{\Sigma}_0$ rather than the ESKF posterior. The ESKF covariance is not propagated into the guard construction. Thus this is a fixed covariance proxy, not a proved prediction-error covariance about $\mathbf{\Phi}(\hat{\mathbf{x}}_{k|k})$. If a tightened row is empty, the implementation collapses it to the midpoint of its untightened interval. This remains an affine nominal constraint but has no chance-constraint interpretation.

Only interior shooting nodes carry these guard rows; neither intermediate sampling points nor the terminal state are constrained by them. Hence the margins do not provide a horizon-wide or continuous-time safety claim. They give $9(N_h - 1)$ two-sided general rows alongside the input box. The learned $\mathbf{\Sigma}_0$ is used here as a fixed guard initialization; its statistical role in ABS-PWM remains the initial latent-belief covariance.

E.5 Exact residual identity and the covariance approximation

Suppose a locally input-equivalent discrepancy acts during the previous sampling interval:

$$\mathbf{z}_k = f_\Delta(\mathbf{z}_{k-1}, \mathbf{u}^{\mathrm{ap}}_{k-1} + \mathbf{d}_k) + \mathbf{w}_{k-1}. \tag{116}$$

Representing actual mismatch in this form is a modeling assumption. In particular, a multiplicative mixer error need not correspond to a constant additive wrench. Define the raw-lift error $\mathbf{n}_k := \mathbf{z}^{\mathrm{m}}_k - \mathbf{z}_k$ and $\mathbf{F}^d_{k-1} := \mathbf{F}(\mathbf{u}^{\mathrm{ap}}_{k-1} + \mathbf{d}_k)$. Bilinearity gives the exact identity

$$\mathbf{e}_k = \mathbf{M}_k \mathbf{d}_k + \mathbf{P}\big(\mathbf{w}_{k-1} + \mathbf{n}_k - \mathbf{F}^d_{k-1} \mathbf{n}_{k-1}\big). \tag{117}$$

Indeed, replacing $\mathbf{z}_{k-1}$ by $\mathbf{z}_{k-1} + \mathbf{n}_{k-1}$ in the nominal predictor adds $\mathbf{F}(\mathbf{u}^{\mathrm{ap}}_{k-1})\mathbf{n}_{k-1}$; the same replacement in $\mathbf{G}_\Delta(\mathbf{z}_{k-1})\mathbf{d}_k$ adds $\sum_i d_{k,i} \mathbf{B}_{\Delta i} \mathbf{n}_{k-1}$. Combining these terms yields (117). Thus the

noise-propagation Jacobian includes the unknown discrepancy. The weight uses its previous stored estimate through $\widehat{\mathbf{F}}_{k-1}$, avoiding an implicit estimation loop.

If coefficients are held fixed and the transition noise and two lift errors are mutually independent and zero mean, with $\mathrm{Cov}(\mathbf{n}_j) \simeq \mathbf{J}_j \boldsymbol{\Sigma}_v \mathbf{J}_j^\top$, then

$$\begin{aligned}\mathrm{Cov}(\mathbf{e}_k - \mathbf{M}_k \mathbf{d}_k) \simeq \mathbf{P}\boldsymbol{\Sigma}_w \mathbf{P}^\top + (\mathbf{P}\mathbf{J}_k)\boldsymbol{\Sigma}_v(\mathbf{P}\mathbf{J}_k)^\top \\ + (\mathbf{P}\mathbf{F}_{k-1}^d \mathbf{J}_{k-1})\boldsymbol{\Sigma}_v(\mathbf{P}\mathbf{F}_{k-1}^d \mathbf{J}_{k-1})^\top.\end{aligned} \tag{118}$$

Substituting $\widehat{\mathbf{F}}_{k-1}$ and adding $\mathbf{E}_k \succ 0$ gives (31). These are additional local assumptions; the observation model alone does not establish them.

To see the limitation, write $h := \widetilde{\boldsymbol{\Phi}}_{\mathrm{ext}}$. The observation equation instead implies

$$\mathbf{n}_k = \boldsymbol{\Phi}(h(\mathbf{z}_k) + \mathbf{v}_k) - \mathbf{z}_k \simeq \underbrace{\boldsymbol{\Phi}(h(\mathbf{z}_k)) - \mathbf{z}_k}_{\mathbf{b}(\mathbf{z}_k)} + \mathbf{J}_k \mathbf{v}_k. \tag{119}$$

The reconstruction term vanishes on consistent physical lifts, but ambient latent states with Gaussian transition noise need not remain on that manifold. It can contribute bias, variance, and correlation with transition noise. For example, with an exact deterministic previous lift, zero discrepancy, and $\mathbf{z}^- := f_\Delta(\mathbf{z}_{k-1}, \mathbf{u}_{k-1}^{\mathrm{ap}})$, let $\mathbf{J}^- := D\boldsymbol{\Phi}(h(\mathbf{z}^-))$ and $\mathbf{K}^- := \mathbf{J}^- Dh(\mathbf{z}^-)$. First-order expansion gives

$$\mathbf{e}_k \simeq \mathbf{P}\mathbf{b}(\mathbf{z}^-) + \mathbf{P}\mathbf{K}^- \mathbf{w}_{k-1} + \mathbf{P}\mathbf{J}^- \mathbf{v}_k. \tag{120}$$

The process contribution is then $\mathbf{P}\mathbf{K}^- \boldsymbol{\Sigma}_w (\mathbf{K}^-)^\top \mathbf{P}^\top$, which need not equal $\mathbf{P}\boldsymbol{\Sigma}_w\mathbf{P}^\top$. This is why (31) is used as a covariance proxy rather than an exact HMM residual covariance. Measured regressors and adjacent residuals are also correlated with their noise. The positive-definite floor ensures an invertible weight, not statistical calibration.

E.6 Analytic measurement-to-lift Jacobian

The raw measurement supplied to the observer has coordinates $\mathbf{y} = (\mathbf{s}, \mathbf{v}, \mathrm{vec}\,\mathbf{R}, \boldsymbol{\omega}) \in \mathbb{R}^{18}$, with column-major vectorization and no additional attitude projection in the lift. For increments $(\delta\mathbf{s}, \delta\mathbf{v}, \delta\mathbf{R}, \delta\boldsymbol{\omega})$, the base-observable differentials are

$$\begin{aligned}\delta\mathbf{p}_0 &= \mathbf{R}^\top \delta\mathbf{s} + (\delta\mathbf{R})^\top \mathbf{s}, \\ \delta\mathbb{Y}_0 &= \mathbf{R}^\top \delta\mathbf{v} + (\delta\mathbf{R})^\top \mathbf{v}, \qquad \delta\mathbb{Z}_0 = \mathrm{vec}(\delta\mathbf{R}).\end{aligned} \tag{121}$$

For a monomial $\mu_a(\boldsymbol{\omega}) = \prod_i \omega_i^{a_i}$,

$$\frac{\partial \mu_a}{\partial \omega_i} = \begin{cases} a_i \omega_i^{a_i - 1} \displaystyle\prod_{j \neq i} \omega_j^{a_j}, & a_i > 0, \\ 0, & a_i = 0. \end{cases} \tag{122}$$

Use the product rule $\delta(\boldsymbol{\mu}_d \otimes \mathbf{p}_0) = (\delta\boldsymbol{\mu}_d) \otimes \mathbf{p}_0 + \boldsymbol{\mu}_d \otimes \delta\mathbf{p}_0$, and similarly for the other base-observable blocks. Stacking these derivatives in the lift's monomial order and dropping the constant row yields $\mathbf{J}_k \in \mathbb{R}^{N \times 18}$. This differentiates all coordinates used in $\mathbf{T}_k$, not only the low-order read block. The learned ambient $\boldsymbol{\Sigma}_v \in \mathbb{R}^{18 \times 18}$ must use this same measurement ordering; the physical tangent covariance $\boldsymbol{\Sigma}_\nu \in \mathbb{R}^{12 \times 12}$ cannot be substituted directly.

E.7 Normalized regression and its statistical interpretation

For fixed $\mathbf{M}_k$ and $\widehat{\mathbf{S}}_k \succ 0$, the implemented estimate minimizes

$$\min_{\mathbf{d} \in \mathbb{R}^4} (\mathbf{e}_k - \mathbf{M}_k \mathbf{d})^\top \mathbf{W}_k (\mathbf{e}_k - \mathbf{M}_k \mathbf{d}) + \lambda \|\mathbf{d}\|_2^2. \tag{123}$$

Multiplication by $c_k > 0$ leaves its minimizer unchanged and gives the unnormalized precision $\widehat{\mathbf{S}}_k^{-1}$ with ridge $\boldsymbol{\Lambda}_k := c_k \lambda \mathbf{I}_4$.

Proposition 3 (Surrogate regression properties)**.** *For $\lambda > 0$, (123) has the unique minimizer (32). Under the stipulated model $\mathbf{e}_k = \mathbf{M}_k \mathbf{d}_k + \boldsymbol{\epsilon}_k$, with fixed regressor, $\boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0}, \widehat{\mathbf{S}}_k)$, and independent prior $\mathbf{d}_k \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Lambda}_k^{-1})$, this minimizer is the posterior mean and MAP estimate, with posterior covariance*

$$(\mathbf{M}_k^\top \widehat{\mathbf{S}}_k^{-1} \mathbf{M}_k + \boldsymbol{\Lambda}_k)^{-1}. \tag{124}$$

For $\lambda = 0$ and full column rank of $\mathbf{M}_k$, the unregularized estimator is unbiased and has minimum covariance among linear unbiased estimators under the same fixed-regressor noise model, without a prior.

Proof. The Hessian in (123) is $2(\mathbf{M}_k^\top \mathbf{W}_k \mathbf{M}_k + \lambda \mathbf{I}_4) \succ 0$. Its normal equations give (32). Completing the square in the Gaussian likelihood multiplied by the prior gives (124) and the stated mean. For the unregularized case, suppress k and let $\mathbf{K} = (\mathbf{M}^\top \widehat{\mathbf{S}}^{-1} \mathbf{M})^{-1} \mathbf{M}^\top \widehat{\mathbf{S}}^{-1}$. Then $\mathbf{K}\mathbf{M} = \mathbf{I}_4$. Any other linear unbiased gain is $\mathbf{K} + \mathbf{D}$ with $\mathbf{D}\mathbf{M} = \mathbf{0}$; its covariance exceeds that of $\mathbf{K}\mathbf{e}$ by $\mathbf{D}\widehat{\mathbf{S}}\mathbf{D}^\top \succeq 0$, since the cross terms vanish. □

These are fixed-regressor Gaussian regression properties [46]. In the controller the regressor and weight depend on noisy measurements, and $\widehat{\mathbf{S}}_k$ is a proxy. The proposition therefore does not assert exact Bayesian inference or unbiased estimation for the physical plant. Its posterior covariance is not the covariance of the subsequently smoothed and projected estimate.

Regularization and smoothing bias. Let $\mathbf{H}_k^d := \mathbf{M}_k^\top \mathbf{W}_k \mathbf{M}_k$ and suppose the fixed-regressor noise is zero mean and $\mathbf{d}$ is deterministic. Then

$$\mathbb{E}[\widetilde{\mathbf{d}}_k] - \mathbf{d} = -\lambda (\mathbf{H}_k^d + \lambda \mathbf{I}_4)^{-1} \mathbf{d}. \tag{125}$$

Along an eigenvector with eigenvalue γ, the mean estimate is shrunk by $\gamma/(\gamma + \lambda)$. This describes an information-matrix eigendirection, not necessarily a single physical wrench coordinate. Normalization changes the effective ridge relative to the unnormalized covariance, so a small numerical λ does not establish negligible bias. All normalization and regularization conventions depend on the specified coordinate units and are held fixed during comparisons. Multiplying the entire covariance proxy by a positive scalar leaves $\mathbf{W}_k$ unchanged. Thus the normalized observer uses relative uncertainty and correlations, rather than responding to a uniform increase in the absolute uncertainty scale.

For constant $\mathbf{d}$, inactive projection, and $\mathbf{b}_k^d := \mathbb{E}[\widetilde{\mathbf{d}}_k] - \mathbf{d}$, smoothing gives

$$\mathbb{E}[\widehat{\mathbf{d}}_k - \mathbf{d}] = (1 - \alpha)\mathbb{E}[\widehat{\mathbf{d}}_{k-1} - \mathbf{d}] + \alpha \mathbf{b}_k^d. \tag{126}$$

Temporal independence is unnecessary for this expectation identity. The homogeneous contraction follows only if $\mathbf{b}_k^d = \mathbf{0}$; a persistent bias is not removed by smoothing. The gain α is a prescribed design parameter. In particular, a ratio of diagonal entries of $\mathbf{\Sigma}_w$ and $\mathbf{\Sigma}_0$ is not generally a steady-state Kalman gain.

E.8 Small-matrix implementation and execution

Cache $\mathbf{P}\mathbf{\Sigma}_w\mathbf{P}^\top$ and form only the selected rows of $\widehat{\mathbf{F}}_{k-1}$ before multiplying by $\mathbf{J}_{k-1}$. After symmetrization and the covariance floor, factor $\widehat{\mathbf{S}}_k = \mathbf{C}_k \mathbf{C}_k^\top$ with a lower Cholesky factor. Solve

$$\begin{aligned} \widehat{\mathbf{S}}_k \mathbf{X}_k = \mathbf{M}_k, \qquad \widehat{\mathbf{S}}_k \mathbf{q}_k = \mathbf{e}_k, \qquad \mathbf{C}_k \mathbf{V}_k = \mathbf{I}_{18}, \\ c_k = \max_a \|[\mathbf{V}_k]_{:,a}\|_2^2, \\ \left(c_k^{-1} \mathbf{M}_k^\top \mathbf{X}_k + \lambda \mathbf{I}_4\right) \widetilde{\mathbf{d}}_k = c_k^{-1} \mathbf{M}_k^\top \mathbf{q}_k. \end{aligned} \tag{127}$$

The column norms are correct because $\widehat{\mathbf{S}}_k^{-1} = \mathbf{C}_k^{-\top} \mathbf{C}_k^{-1}$. This avoids forming the dense covariance inverse, although the inverse triangular factor is computed to obtain the normalization. The observer requires an eighteen-dimensional factorization and a four-dimensional regression solve, with no propagated $N \times N$ belief covariance and no extra QP variables. A diagonal-proxy variant replaces $\widehat{\mathbf{S}}_k$ by its diagonal before normalization; a process-only comparison uses the diagonal of $\mathbf{P}\mathbf{\Sigma}_w\mathbf{P}^\top$. At the first sample, initialize the ESKF, store the raw lift and its Jacobian, and use a zero wrench estimate. At each subsequent sample, predict and update the ESKF to obtain $\hat{\mathbf{x}}_{k|k}$. Form the observer residual using the previous applied command, construct the weight using the previous stored estimate, solve the regression, and smooth and project once. Solve the nominal QP from $\mathbf{\Phi}(\hat{\mathbf{x}}_{k|k})$ and apply (5) once, storing the actual command for the next prediction. Guard and observer quantities are fixed before their respective solves. Numerical fallback events and infeasible QPs require separately specified recovery policies; the regression and QP propositions do not establish their closed-loop behavior.

E.9 Scope of the uncertainty and control claims

The physical-state ESKF filters the MPC initial condition using its nominal model and calibrated sensor covariance. The learned ABS-PWM covariances have separate roles: $\mathbf{\Sigma}_w$ and $\mathbf{\Sigma}_v$ determine the observer's precision metric, while $\mathbf{\Sigma}_0$ and $\mathbf{\Sigma}_w$ initialize and propagate the precomputed nominal guard proxy. Neither mechanism uses the ESKF posterior covariance. The covariances are not claimed to be calibrated physical error distributions merely because they were learned by ABS-PWM.

Algorithm 2: BUC-MPC with ESKF state estimation and covariance-weighted wrench compensation

Input: ABS-PWM operators and covariances $\mathbf{\Sigma}_w, \mathbf{\Sigma}_v, \mathbf{\Sigma}_0$; reference trajectory; sampling interval Δt, shooting interval $\delta = n_s \Delta t$, and horizon N_h; MPC weights and bounds; observer parameters $\lambda > 0$, $0 < \alpha \leq 1$, $d_{\max}$, and $\mathbf{E}_k \succ 0$; configured ESKF and bounded recovery policy.

Output: Applied wrench commands $\mathbf{u}_k^{\mathrm{ap}}$.

Precompute $\mathbf{A}_d, \mathbf{\Gamma}_\delta$ by (27); choose $\mathbf{Q}_f \succeq 0$ as in § E.3
Precompute the guard covariance proxies and marginal deviations $\boldsymbol{\sigma}_l$ by (114), $l \in [1 : N_h - 1]$
Cache $\mathbf{P}\mathbf{\Sigma}_w\mathbf{P}^\top$ and set $\hat{\mathbf{d}}_0 \leftarrow \mathbf{0}$
for $k = 0, 1, 2, \ldots$ **do**
    Acquire measurement $\mathbf{y}_k$
    if $k = 0$ **then**
        Initialize the ESKF from $\mathbf{y}_0$ to obtain $\hat{\mathbf{x}}_{0|0}$, as in § E.1
    else
        Predict the ESKF using $\mathbf{u}_{k-1}^{\mathrm{ap}}$ and update with $\mathbf{y}_k$ to obtain $\hat{\mathbf{x}}_{k|k}$
    Form the raw lift $\mathbf{z}_k^{\mathrm{m}} \leftarrow \mathbf{\Phi}(\mathbf{y}_k)$ and its Jacobian $\mathbf{J}_k \leftarrow D\mathbf{\Phi}(\mathbf{y}_k)$
    if $k \geq 1$ **then**
        Form the raw-measurement residual $\mathbf{e}_k$ and regressor $\mathbf{M}_k$ by (30), using $\mathbf{z}_{k-1}^{\mathrm{m}}$ and $\mathbf{u}_{k-1}^{\mathrm{ap}}$
        Form $\widehat{\mathbf{F}}_{k-1}$ and the covariance proxy $\widehat{\mathbf{S}}_k$ by (31), using the stored estimate $\hat{\mathbf{d}}_{k-1}$
        Normalize the precision and solve for $\widetilde{\mathbf{d}}_k$ by (32), using the small-matrix solves in (127)
        Smooth and project once: $\hat{\mathbf{d}}_k \leftarrow \mathrm{proj}_{\mathcal{D}}\big(\hat{\mathbf{d}}_{k-1} + \alpha(\widetilde{\mathbf{d}}_k - \hat{\mathbf{d}}_{k-1})\big)$
    Set $\bar{\mathbf{z}}_{0|k} \leftarrow \mathbf{\Phi}(\hat{\mathbf{x}}_{k|k})$
    Fix guide states $\mathbf{z}_{l|k}^{\mathrm{g}}$ from the reference or shifted previous predictions, as in § E.2
    Form $\mathbf{B}_{d,l|k}$ by (27), $l \in [0 : N_h - 1]$, and the guard bounds by (113)–(115)
    Condense the nominal predictions and solve (29) for $\mathbf{U}_k^\star$
    if *the QP returns a usable feasible solution* **then**
        Set $\mathbf{u}_k^{\mathrm{ap}} \leftarrow \mathrm{clip}_{[\mathbf{u}_{\min}, \mathbf{u}_{\max}]}(\mathbf{u}_{0|k}^\star - \hat{\mathbf{d}}_k)$
        Cache the optimal predictions for subsequent guide construction
    else
        Obtain $\mathbf{u}_k^{\mathrm{ap}} \in [\mathbf{u}_{\min}, \mathbf{u}_{\max}]$ from the specified recovery policy
    Apply $\mathbf{u}_k^{\mathrm{ap}}$ for Δt
    Store $\mathbf{z}_k^{\mathrm{m}}$, $\mathbf{J}_k$, $\hat{\mathbf{d}}_k$, and the actual applied command $\mathbf{u}_k^{\mathrm{ap}}$ for the next sample

The residual predictor uses the command actually sent, without adding an estimated discrepancy. The discrepancy enters only the covariance Jacobian and the final command correction. If the true additive discrepancy equals the estimate and command clipping is inactive, the effective first wrench matches the QP's planned wrench. State prediction can nevertheless differ because the learned transition and fixed-guide shooting model approximate the plant. Post-QP correction enforces the command box but supplies no independent state-constraint guarantee.

The convexity and regression propositions establish well-posed computations. They do not establish recursive feasibility, exact offset-free regulation, closed-loop stability, or safety between sampling points. The benefit of covariance weighting is an empirical question: an observer-off comparison isolates compensation, whereas covariance-independent, process-only, diagonal-proxy, and full-proxy comparisons distinguish the roles of learned weighting and correlations. Such comparisons must preserve the coordinate scales, regularization convention, and other controller settings.

F SUPPLEMENTARY MATERIAL FOR § 6

F.1 PRELIMINARY INFO FOR EXPERIMENTS

Truncation degree for ABS-PWM. We use $K = 1$ throughout unless otherwise stated, with justification provided in a later subsection.

Description of an unsuccessful flight. A flight is considered unsuccessful if it triggers any flight-envelope termination: physical contact in MuJoCo, altitude below $0.15\,\mathrm{m}$, tilt beyond 120°, position beyond $30\,\mathrm{m}$, speed above $15\,\mathrm{m\,s^{-1}}$, any body-rate component above $20\,\mathrm{rad\,s^{-1}}$, or a non-finite state.

Disturbances in MuJoCo. The same disturbance model is active during both training-data collection and closed-loop evaluation. Disturbances are applied as external forces and torques at every 2 ms MuJoCo physics substep and are logged separately from the applied control wrench. This distinction is intentional: the applied control wrench represents the actuation delivered by the rotors, whereas the external disturbance wrench represents unmodelled dynamics that the residual model must capture. Let $v_{\mathrm{rel},k} = v_k - (w + \gamma_k)$ denote the air-relative velocity at physics step k. The external force in the world frame and the external torque in the body frame are

$$F_k^{\mathrm{dist}} = -D_{\mathrm{lin}} \odot v_{\mathrm{rel},k} - D_{\mathrm{quad}} \odot (|v_{\mathrm{rel},k}| \odot v_{\mathrm{rel},k}) + b_k \sigma_F \xi_k, \tag{128}$$

$$T_{\mathrm{body},k}^{\mathrm{dist}} = -D_{\mathrm{ang}} \odot \omega_k + b_k \sigma_\tau \zeta_k, \tag{129}$$

where $\odot$ denotes element-wise multiplication and $\xi_k, \zeta_k \sim \mathcal{N}(0, I_3)$. The disturbance components are defined as follows:

- $D_{\mathrm{lin}}, D_{\mathrm{quad}} \in \mathbb{R}^3$ are the per-world-axis linear and quadratic drag coefficients, respectively, while $D_{\mathrm{ang}} \in \mathbb{R}^3$ contains the per-body-axis angular-damping coefficients. All three vectors are sampled independently.
- $w \in \mathbb{R}^3$ is a constant world-frame wind velocity.
- $\gamma_k \in \mathbb{R}^3$ is an Ornstein–Uhlenbeck gust process about the constant wind. It is discretized exactly at the physics time step as $\gamma_{k+1} = a\gamma_k + \sqrt{1-a^2}\,\sigma_\gamma \eta_k$, where $a = \exp(-\Delta t/\tau_\gamma)$ and $\eta_k \sim \mathcal{N}(0, I_3)$.
- Impulsive disturbances are generated by $b_k \sim \mathrm{Bernoulli}(\lambda \Delta t)$, which is the discrete-time form of a Poisson process with rate λ. When $b_k = 1$, an independent Gaussian force and torque are applied for one physics substep.

MuJoCo's medium density and viscosity, which generate additional aerodynamic forces on the geometric primitives attached to the vehicle body, are also randomized. Distinct from external disturbances, the actuator model includes first-order rotor dynamics with time constant τ_m, per-rotor variations in thrust and drag-torque coefficients, and multiplicative actuator noise.

Observation noise. To emulate sensor noise, the measured state is

$$\mathbf{y}_k = \mathbf{x}_k \boxplus \boldsymbol{\nu}_k, \qquad \boldsymbol{\nu}_k = \begin{bmatrix} \boldsymbol{\nu}_{p,k}^\top & \boldsymbol{\nu}_{v,k}^\top & \boldsymbol{\nu}_{R,k}^\top & \boldsymbol{\nu}_{\omega,k}^\top \end{bmatrix}^\top \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\nu) \quad \text{i.i.d.},$$

where $\boldsymbol{\nu}_{p,k}$, $\boldsymbol{\nu}_{v,k}$, and $\boldsymbol{\nu}_{\omega,k}$ are added to position, velocity, and body rate, respectively, while attitude is perturbed on $\mathrm{SO}(3)$ as $\mathbf{R}_k^y = \mathbf{R}_k \exp(\widehat{\boldsymbol{\nu}}_{R,k})$. The four isotropic noise components have standard deviations

$$\sigma_p = 10\,\mathrm{mm}, \qquad \sigma_v = 0.02\,\mathrm{m\,s^{-1}}, \qquad \sigma_R = 0.0087\,\mathrm{rad}\ (0.5^\circ), \qquad \sigma_\omega = 0.010\,\mathrm{rad\,s^{-1}}.$$

Neither the learning procedure nor the controller observes the true state $\mathbf{x}_k$ or the sampled realization $\boldsymbol{\nu}_k$; both operate solely on the noisy measurement $\mathbf{y}_k$. Observation noise is applied at every time step, both during training-data generation and throughout the closed-loop experiments.

In-distribution (ID) and out-of-distribution (OOD) randomization. We randomize the physical and environmental domain in MuJoCo for two conditions: ID and OOD as highlighted in Table 4. The measurement noise is fixed in both conditions.

Table 4: Physical and environmental randomization for the ID and OOD settings. Unless otherwise stated, every scalar is sampled uniformly over the reported interval, and vector components and per-rotor quantities are sampled independently.

Quantity	Type	ID range	Implemented OOD setting
Airframe mass	relative	$\pm 10\%$ of nominal	$\pm 20\%$ of nominal
Payload mass	absolute	$[0, 0.15]$ kg	$[-0.175, 0.525]$ kg; a negative draw is subtracted from the randomized airframe mass
Inertia, each principal moment	relative	$\pm 10\%$ of nominal	$\pm 20\%$ of nominal, subject to positive-definiteness and the rigid-body triangle inequality
Inertia principal-axis rotation	absolute	$[-0.05, 0.05]$ rad per axis	$[-0.10, 0.10]$ rad per axis
Centre-of-mass offset	absolute	$[-0.010, 0.010]$ m per axis	$[-0.020, 0.020]$ m per axis
Thrust coefficient, each rotor	relative	$\pm 10\%$	$\pm 20\%$
Drag-torque coefficient, each rotor	relative	$\pm 10\%$	$\pm 20\%$
Motor time constant τ_m	absolute	$[0.015, 0.045]$ s	$[-0.020, 0.100]$ s; draws with $\tau_m \leq 10^{-9}$ s produce instantaneous rotor response
Linear drag D_{lin}	absolute	$[0.02, 0.20]$ N/(m/s) per axis	$[-0.07, 0.29]$ N/(m/s) per axis; negative values act as linear anti-damping
Quadratic drag D_{quad}	absolute	$[0.01, 0.12]$ N/(m/s)^2 per axis	$[-0.045, 0.175]$ N/(m/s)^2 per axis; negative values act as quadratic anti-damping
Angular drag D_{ang}	absolute	$[0.5, 4.0] \times 10^{-3}$ N m/(rad/s)	$[-1.25, 5.75] \times 10^{-3}$ N m/(rad/s); negative values act as rotational anti-damping
Constant wind w	absolute	$[-1, 1]$ m/s per axis	$[-5, 5]$ m/s per axis
Gust standard deviation σ_γ	absolute	$[0, 0.5]$ m/s	$[-0.75, 2.25]$ m/s; the Ornstein–Uhlenbeck gust is disabled when $\sigma_\gamma \leq 0$
Gust correlation time τ_γ	absolute	$[0.4, 2.0]$ s	$[-0.4, 2.8]$ s; the Ornstein–Uhlenbeck gust is disabled when $\tau_\gamma \leq 10^{-6}$ s
Impulse rate λ	absolute	$[0, 0.3]$ s^{-1}	$[-0.5, 1.5]$ s^{-1}; impulsive disturbances are disabled when $\lambda \leq 0$
Impulse-force scale σ_F	absolute	$[0, 1.0]$ N	$[-1.5, 4.5]$ N; because this scale multiplies a zero-mean Gaussian variable, its sign does not change the force distribution
Impulse-torque scale σ_τ	absolute	$[0, 0.02]$ N m	$[-0.03, 0.09]$ N m; because this scale multiplies a zero-mean Gaussian variable, its sign does not change the torque distribution
Air density and viscosity	relative	$\pm 10\%$	$\pm 20\%$
Actuator-noise scale	absolute	$[0, 0.01]$	$[-0.005, 0.015]$; actuator noise is applied only when $\sigma_a > 0$

Skydio X2 nominal parameters. The nominal parameters of Skydio X2 in MuJoCo are

$$m = 1.325\,\text{kg}, \qquad g = 9.81\,\text{m/s}^2, \qquad \mathbf{J} = \begin{bmatrix} 0.03665 & 0 & -0.0021 \\ 0 & 0.02541 & 0 \\ -0.0021 & 0 & 0.0605 \end{bmatrix} \text{kg} \cdot \text{m}^2.$$

The actuation limits $f \in [10, 42]$ N, $|\tau_1| \leq 0.600$, $|\tau_2| \leq 0.467$, and $|\tau_3| \leq 0.067$.

F.2 Training data collection

Because quadrotors are inherently unstable, a nominal stabilizing controller is required for stable training data collection. We therefore employ the geometric controller proposed in [45] for the generation

of training data with geometrically consistent references generated using [56]. Each trajectory lasts for a time horizon of 5 s and collected at 100 Hz sampling ($\Delta t = 0.01$ s, 500 control steps). We collect 1.5 min of flight data across nine different reference families for training as highlighted in Table 5 in ID setting. Except for recovery families, the initial position and velocity offsets from the reference are sampled uniformly from the boxes $\Delta \mathbf{s}_x \in \pm(0.25, 0.25, 0.20)$ m and $\Delta \mathbf{v}_x \in \pm(0.30, 0.30, 0.20)$ m/s, respectively. The initial attitude is perturbed by a random-axis rotation with magnitude in $[0, 0.15]$ rad, and the body-rate offset is sampled uniformly from $\Delta\boldsymbol{\omega} \in \pm(0.30, 0.30, 0.30)$ rad/s. The recovery families use wider bounds: $\Delta \mathbf{s}_x \in \pm(1.50, 1.50, 0.90)$ m, $\Delta \mathbf{v}_x \in \pm(1.80, 1.80, 1.00)$ m/s, an attitude perturbation of up to 0.60 rad, and $\Delta\boldsymbol{\omega} \in \pm(2.50, 2.50, 2.00)$ rad/s. The recorded control input $\mathbf{u}_k$ comprises the nominal controller action and a bounded probing excitation, after enforcing the admissible wrench limits. Persistent excitation can be promoted during data collection by augmenting the nominal controller action with a bounded, spectrally rich probing signal. Excitation can be assessed over finite windows using the smallest eigenvalue of an input Gram matrix, with insufficiently excited trajectories rejected or recollected.

Table 5: Training data: trajectory families and the ranges from which each trajectory instance is drawn. Continuous parameters are sampled uniformly over the stated intervals.

Family	Trajectories	Instance parameters
hover	2	Hover point: $(x, y) \in [-2, 2]^2$ and $z \in [0.8, 2.5]$ m.
recovery	2	Hover point: $(x, y) \in [-2, 2]^2$ and $z \in [1.2, 2.5]$ m; wide initial-condition block.
disturbance recovery	2	Hover point: $(x, y) \in [-2, 2]^2$ and $z \in [1.2, 2.5]$ m; wide initial-condition block and boosted disturbances with $\lambda = 1.2\,\mathrm{s}^{-1}$, $\sigma_F = 2.0$ N, $\sigma_\tau = 0.05$ N m, and $\sigma_\gamma \geq 0.8$ m/s..
circle	2	Radius $[0.5, 2.0]$ m, angular rate $[0.4, 1.8]$ rad/s, direction ± 1, and random phase.
helix	2	Circular motion with radius $[0.5, 1.8]$ m, angular rate $[0.4, 1.6]$ rad/s, upward climb rate $[0.25, 0.40]$ m/s (1.25–2.0 m over an episode), and a vertical sinusoid with amplitude $[0, 0.15]$ m and frequency $[0.05, 0.35]$ Hz.
lemniscate	2	Bernoulli figure-eight with radius $[0.6, 2.0]$ m, angular rate $[0.4, 1.4]$ rad/s, and plane tilt about the x-axis in $[-0.35, 0.35]$ rad.
weave	2	Three-dimensional Lissajous trajectory $p = c + A \odot \sin(2\pi F t + \phi)$, with $A_x, A_y \in [0.90, 1.50]$ m, $A_z \in [0.30, 0.60]$ m, and $F_i \in [0.20, 0.34]$ Hz sampled independently for $i \in \{x, y, z\}$; random phases are sampled independently for each axis.
point_to_point	2	2–4 minimum-snap segments, rest-to-rest at every knot; hops of $[0.8, 2.5]$ m per horizontal axis and $[0, 1.0]$ m vertically, with random signs; segment durations are jittered by $\pm 30\%$ about equal spacing.
aggressive	2	4–7 minimum-snap segments with pass-through (Catmull–Rom) knot velocities; horizontal hops of $[1.5, 3.5]$ m and vertical hops of $[0, 1.2]$ m. Short segments and long hops generate hard acceleration and braking manoeuvres.
Dataset summary		
Total trajectories (M)	18	Two trajectories from each of the nine families.
Duration per trajectory	5 s	Closed-loop flight duration.
Control and logging rate	100 Hz	Sampling interval $\Delta t = 0.01$ s.
MuJoCo physics integration step	2 ms	Five simulation substeps per control step.
State samples per trajectory	501	Includes the initial and terminal states.
Transitions per trajectory	500	Each transition is $(\mathbf{y}_k, \mathbf{u}_k, \mathbf{y}_{k+1})$.
Total state samples	9,018	18×501 state observations.
Total transitions	9,000	18×500 state–input transitions.
Total flight time	90 s	18×5 s of simulated flight data.

F.3 Test data for Open-loop Evaluation

We evaluate the learned models on two test datasets that are separate from the training data. The in-distribution (ID) test dataset consists of 74 trajectories, whereas the out-of-distribution (OOD)

test dataset consists of 174 independently collected episodes generated using the physical and environmental randomization reported in Table 4. Neither test dataset is used for model training.

Both test datasets contain the same nine reference families and use the same trajectory-parameter ranges, time duration, sampling rate, observation-noise model, and family-dependent initial-condition distributions as the training data. Consequently, the ID and OOD test datasets differ in their physical and environmental randomization rather than in their reference trajectories or initial-condition distributions. Each test episode is a 5 s closed-loop flight sampled at 100 Hz and contains 500 state–input transitions. The ID and OOD datasets therefore contain 37,000 and 87,000 transitions, corresponding to 370 s and 870 s of flight data, respectively. Their composition is reported in Table 6.

Table 6: Composition of the separate ID and OOD test datasets. Each trajectory contributes eight one-second prediction windows.

Family	ID test		OOD test	
	Trajectories	Windows	Trajectories	Windows
hover	1	8	11	88
recovery	6	48	21	168
disturbance recovery	6	48	6	48
circle	8	64	33	264
helix	10	80	25	200
lemniscate	16	128	23	184
weave	12	96	27	216
point_to_point	11	88	20	160
aggressive	4	32	8	64
Total	74	592	174	1,392

For open-loop evaluation, eight windows of 100 transitions, corresponding to a prediction horizon of 1 s, are extracted from each test trajectory. The windows begin at $0.5, 1.0, \ldots, 4.0$ s, so consecutive windows overlap by 0.5 s. These windows are evaluation segments extracted from the recorded trajectories and are not additional independently generated trajectories. Each rollout is initialized from the noisy state measurement at the beginning of the window and propagated using the corresponding sequence of recorded commanded inputs $\mathbf{u}_k$. Prediction errors are computed against the latent true state $\mathbf{x}_k$. Models requiring a state history are additionally provided with the recorded context preceding the window boundary.

Selection of the truncation degree for ABS-PWM. We selected truncation degree $K = 1$ because higher truncation degrees substantially increase both training and inference costs. A GECM iteration is approximately $4.1\times$ slower for $K = 2$ and $33.7\times$ slower for $K = 3$ relative to $K = 1$. Moreover, the $K = 3$ open-loop rollout alone requires approximately 22 ms, exceeding the 10 ms computation budget imposed by the 100 Hz sampling rate, even before accounting for state estimation and control optimization. The $K = 1$ model therefore provides the greatest computational margin for real-time operation.

Table 7: Computational cost as a function of the ABS-PWM truncation degree. GECM times are reported as mean $\pm$ sample standard deviation; rollout times (mean) correspond to a 1 s prediction horizon.

K	Lifted dimension	GECM step time [s]	Rollout time [ms]
1	63	2.77 ± 0.81	1.27
2	159	11.33 ± 2.05	5.61
3	319	93.25 ± 25.34	22.02

F.4 Closed-loop Implementation Details

The learned physical parameters and covariances remain fixed throughout evaluation, including under OOD conditions. For BUC-MPC and all benchmarks other than SkyJEPA, we use $n_s = 10$, $\delta = 0.1$ s, and a prediction horizon of 1 s, while retaining the remaining configurations reported by the respective authors. We implement SkyJEPA+MPPI on a GPU in PyTorch using the configuration reported in [8]. The convex QP of BUC-MPC is solved using DAQP solver [59].

State estimation. We use the same ESKF output as the state input to every controller. State estimation is treated as shared system infrastructure rather than part of the learned world model. This prevents differences in measurement filtering from confounding the comparison and ensures that every method is initialized from the same causal, geometrically consistent physical-state estimate. The ESKF uses only nominal rigid-body dynamics and the common sensor-noise model; it does not use the learned dynamics of any evaluated method. This separation is also consistent with physical quadrotor deployments, including KoopNet, where sensor fusion supplies a filtered state estimate before lifting and predictive control [35].

The ESKF uses a 12-dimensional right-error state ordered as position, velocity, attitude error, and body rate. Its nominal predictor integrates the rigid-body dynamics by RK4 using the previously applied wrench, nominal mass, and nominal inertia; drag and mixer corrections are omitted. The per-sample process covariance is

$$\mathbf{Q}_{\text{est}} = \operatorname{diag}\big((10^{-3}\,\text{m})^2\mathbf{I}_3,\ (2\times10^{-2}\,\text{m}\,\text{s}^{-1})^2\mathbf{I}_3,\ (10^{-3}\,\text{rad})^2\mathbf{I}_3,\ (5\times10^{-2}\,\text{rad}\,\text{s}^{-1})^2\mathbf{I}_3\big)\,.$$

The first measurement initializes the nominal state. Subsequent corrections use the Joseph covariance update and an identity approximation for the attitude-error reset Jacobian. The transition Jacobian is generated with CasADi and compiled when a C compiler is available. The filtered state initializes the MPC as $\bar{\mathbf{z}}_{0|k} = \mathbf{\Phi}(\hat{\mathbf{x}}_{k|k})$. The learned ambient covariance $\mathbf{\Sigma}_v$ is used by the wrench observer and is separate from the ESKF.

Guide construction and reference tracking. At the start of each flight, the guide trajectory is initialized from the reference at the shooting nodes. After each successful solve, the predicted lifted states are shifted by one shooting interval, and the final node is initialized from the reference. At the next control step, the first guide state is replaced by the current filtered state, while the remaining guide states are reconstructed and re-lifted. The guide-dependent input matrices $\mathbf{H}(\mathbf{z}^{\text{g}}_{l|k})$ are then evaluated once and held fixed during the QP solve. One guide update and one QP solve are performed per control step, without successive linearization. The guide is shifted by $\delta = 0.1\,\text{s}$ at each $\Delta t = 0.01\,\text{s}$ control step. Tracking targets and feedforward wrenches are evaluated at $k\Delta t + l\delta$ using the differential-flatness reference generator [56].

Table 8: BUC-MPC parameters used for the closed-loop experiments.

Quantity	Setting
Sampling interval	$\Delta t = 0.01\,\text{s}$
Shooting interval	$\delta = 0.1\,\text{s}$, $n_s = 10$
Horizon	$N_h = 10$
Lifted dimension	$N = 63$
Stage-state weight $\mathbf{Q}$	$\operatorname{diag}(1000\mathbf{I}_3,\ 60\mathbf{I}_3,\ 1200\mathbf{I}_9,\ 12\mathbf{I}_3,\ 0.01\mathbf{I}_{45})$
Input weight $\mathbf{R}_u$	$\operatorname{diag}(1, 0.2, 0.2, 0.2)$
Terminal weight $\mathbf{Q}_f$	DARE solution with virtual-input weight $\rho_{\text{term}}\mathbf{I}$, where $\rho_{\text{term}} = 0.1$
Input bounds	$\mathbf{u}_{\min} = (10, -0.60, -0.467, -0.067)^\top$ and $\mathbf{u}_{\max} = (42, 0.60, 0.467, 0.067)^\top$, ordered as $(f, \tau_1, \tau_2, \tau_3)$ in N and N m
State envelopes	$\mathbf{s}_{\max} = (30, 30, 30)\,\text{m}$, $\mathbf{v}_{\max} = (6, 6, 4)\,\text{m}\,\text{s}^{-1}$, and $\boldsymbol{\omega}_{\max} = (8, 8, 5)\,\text{rad}\,\text{s}^{-1}$
Guard multiplier	$\kappa = 1$
Observer ridge	$\lambda = 10^{-6}$
Observer smoothing gain	$\alpha = 0.018384105$
Observer projection radius	$d_{\max} = 3$
Residual covariance regularizer	$\mathbf{E}_k = 10^{-12}\mathbf{I}_{18}$
QP solver	DAQP
Solver tolerances	10^{-7} primal and dual
Solver iteration limit	2000

F.5 CRAZYFLIE 2 WITH BUC-MPC

We also consider BUC-MPC implementation for Crazyflie 2 drone after collecting similar data like Skydio X2 and training ABS-PWM. Qualitatively, OOD tracking is found to be better than Skydio X2.

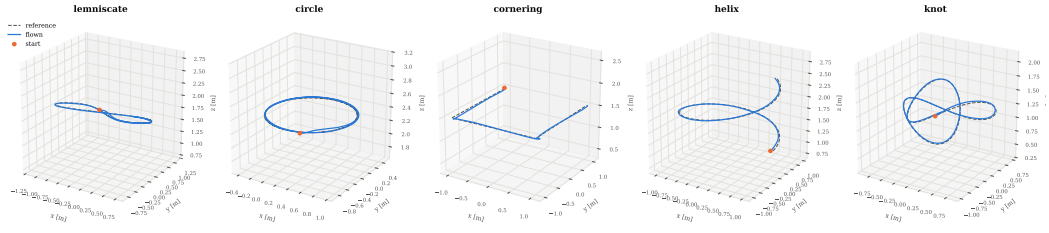


Figure 8: Representative ID trajectory-tracking demonstration with BUC-MPC for the Crazyflie 2.

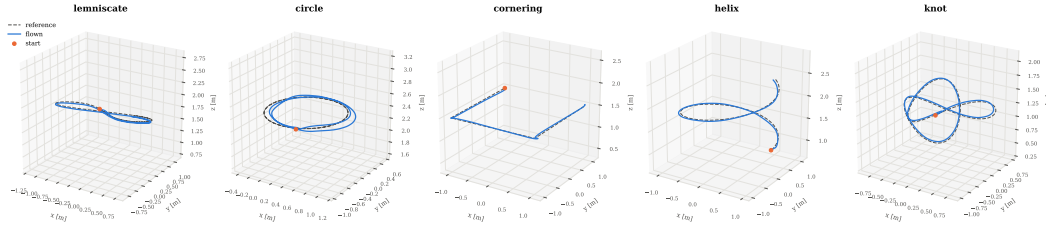


Figure 9: Representative OOD trajectory-tracking demonstration with BUC-MPC for the Crazyflie 2.

F.6 ADDITIONAL RESULTS

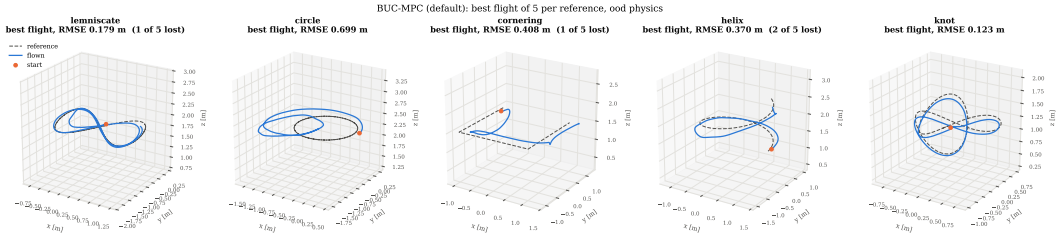


Figure 10: Representative OOD trajectory-tracking demonstration with BUC-MPC for the Skydio X2.

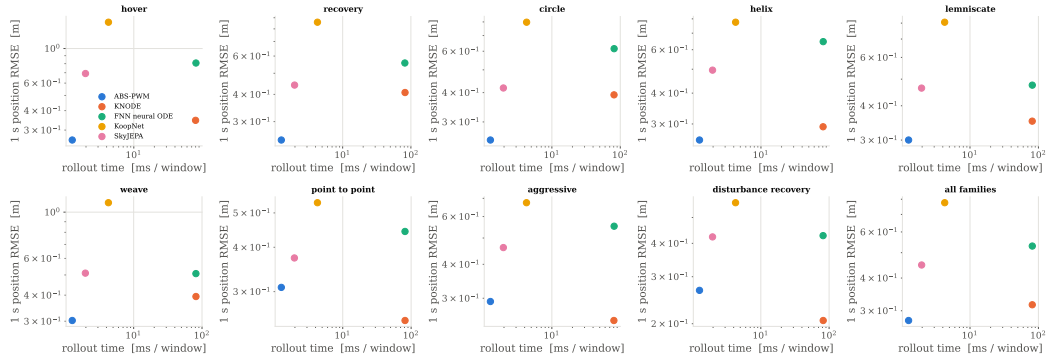


Figure 11: Test trajectory family wise open-loop prediction comparison for Skydio X2 in ID setting

F.7 FAILURE OF KOOPNET BENCHMARK IN CLOSED-LOOP

Under the evaluated configuration, KoopNet completed none of the 50 trials in either ID or OOD conditions, as all flights violated the prescribed flight-envelope constraints. KoopNet's unsuccessful flights are consistent with its larger prediction errors in Table 1, particularly the substantial ID velocity error and elevated position errors under both conditions. Such errors can compromise the predicted

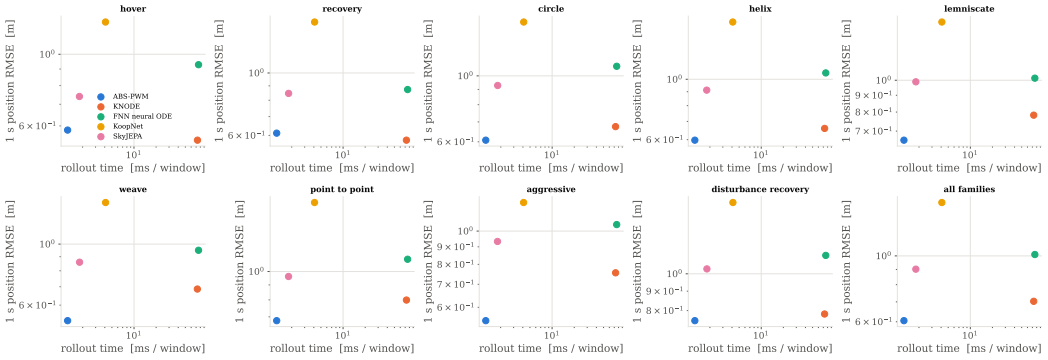


Figure 12: Test trajectory family wise open-loop prediction comparison for Skydio X2 in OOD setting

consequences of candidate control inputs, leading to tracking deviations and eventual violations of the flight-envelope constraints. This provides a plausible explanation for the observed failures, although the aggregate prediction results do not isolate their cause.

G Ablation Study

All ablation studies conducted for Skydio X2 using the same test and train data.

G.1 Does a learned decoder help?

We replace the analytical decoder with a 63–126–18 MLP while keeping the remaining components of AWS-PWM fixed. The analytical decoder provides substantially more accurate position predictions, particularly at short horizons, where the learned decoder exhibits a large initial readout error. The learned decoder improves attitude prediction but produces less accurate body-rate predictions. Overall, the analytical readout is more effective for short-horizon position prediction.

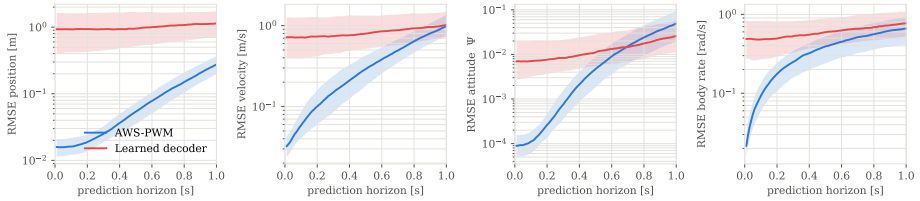


Figure 13: Decoder ablation: Curves show the median cumulative RMSE over the prediction horizon, and shaded regions denote the interquartile range.

G.2 How does analytical lift, decoding, and model-based inference help ?

We replace the analytical lift, analytical decoder, and GECM inference with a bidirectional GRU encoder, an MLP decoder, and gradient-based amortized inference, while retaining the same 63-dimensional structured transition. The amortized model is less accurate across all state components in the. The gap is already present at short horizons and grows throughout the rollout, especially for position and attitude. This shows that the analytical representation and model-based inference are important when fitting the structured dynamics from limited data.

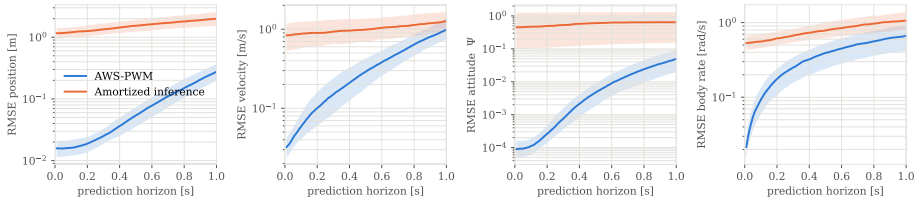


Figure 14: Amortized-inference ablation: Curves show the median cumulative RMSE over the prediction horizon; shaded regions denote the interquartile range.

G.3 How does the structured bilinear transition help?

We replace the analytically constructed lift and bilinear operators with a generic 63-dimensional bilinear model while keeping the latent dimension and training data fixed. AWS-PWM provides markedly more accurate position predictions, with the largest advantage at short horizons where the generic model already incurs an initial-state inference error. As seen in Fig. 15, the generic model yields lower velocity, attitude, and body-rate errors, but its position error remains higher throughout the one-second rollout. The analytical structure therefore improves the data efficiency of the position prediction used for receding-horizon tracking.

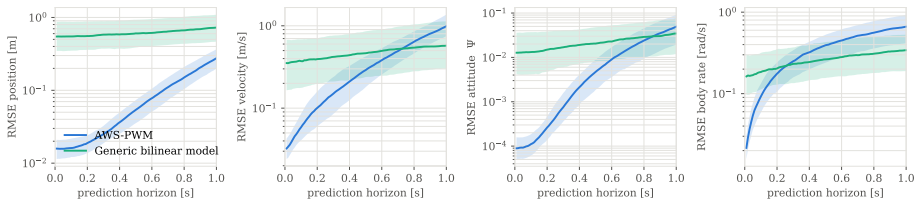


Figure 15: Bilinear-structure ablation: Curves show the median cumulative RMSE over the prediction horizon; shaded regions denote the interquartile range.

G.4 PHYSICAL-PARAMETER INITIALIZATION.

We vary the initial mass, inertia, and damping parameters while keeping the training data, covariance initialization, and GECM settings fixed. All physically feasible initializations converge, including perturbations that scale the initial parameters by factors between 0.5 and 2. Convergence requires 111–124 GECM iterations across the tested initializations, compared with 117 iterations for the default initialization, indicating limited sensitivity in convergence time. Their ID position errors remain close to the default initialization, while OOD errors vary modestly and non-monotonically. Thus, an accurate physical initialization is not required for convergence or ID prediction, although it can affect the identified physical solution and OOD generalization.

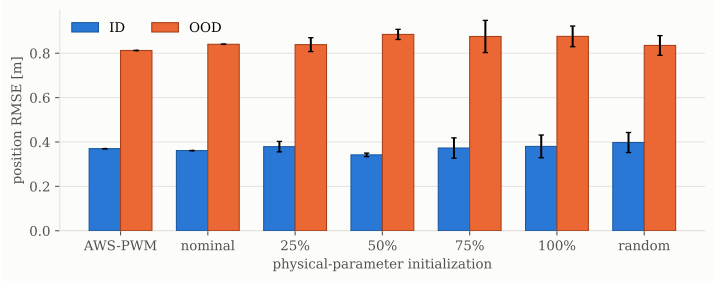


Figure 16: Sensitivity to physical-parameter initialization for the. Bars show the mean position RMSE; error bars show variation across initializations.

G.5 OBSERVATION-NOISE ABLATION.

We increase the observation-noise covariance used to generate the training data while keeping the trajectories and GECM settings fixed. Doubling the covariance does not degrade position prediction: the ID and OOD errors remain close to the baseline across all tested noise levels. All fits converge, requiring 121–123 GECM iterations at the increased noise levels, compared with 117 iterations for AWS-PWM at the original noise level. The corresponding fitting times are approximately 6.2–6.9 minutes, compared with 5.4 minutes for the baseline. Thus, the additional observation noise has little effect on prediction accuracy and only modestly increases convergence time in this range.

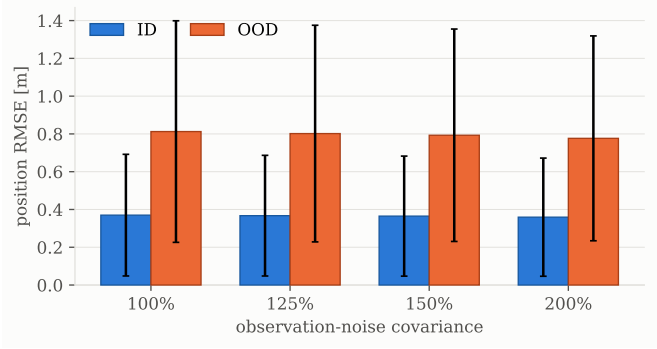
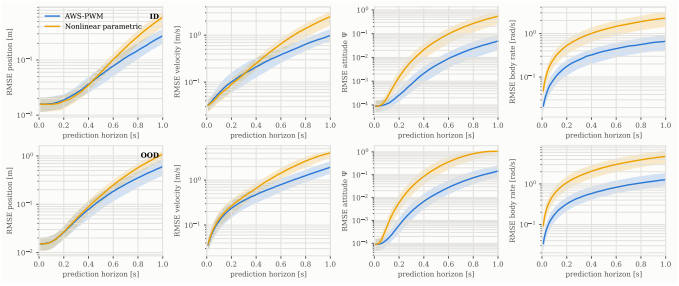


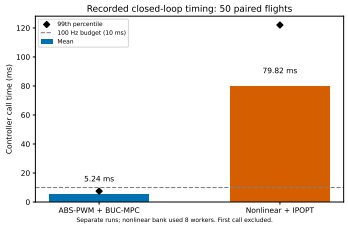
Figure 17: Observation-noise ablation: noise levels denote the total covariance relative to the original data. Bars show the mean position RMSE, and error bars show variation across test trajectories.

G.6 NONLINEAR-PARAMETRIC MODEL ABLATION.

We compare AWS-PWM with a constrained nonlinear rigid-body model fitted directly in the physical state using an error-state EKF. Both models use the same training trajectories. AWS-PWM produces lower open-loop errors for all state components under both ID and OOD conditions, with the gap increasing over the prediction horizon. For closed-loop evaluation, AWS-PWM uses BUC-MPC with a condensed QP solved by DAQP, whereas the nonlinear model uses a direct multiple-shooting nonlinear MPC solved by IPOPT. Both controllers complete all 50 ID trials, but BUC-MPC gives lower position error and requires 5.24 ms per step, compared with 79.82 ms for the nonlinear MPC. The nonlinear fit takes 13.6 minutes, while AWS-PWM converges in 5.4 minutes.



(a) Median open-loop RMSE under ID and OOD conditions. Shaded regions show the interquartile range.



(b) Closed-loop computation time.

Figure 18: Nonlinear-parametric-model ablation. AWS-PWM uses BUC-MPC with DAQP, while the nonlinear physical model uses multiple-shooting nonlinear MPC with IPOPT.

H LIMITATIONS

ABS-PWM uses EKF–RTS smoothing to infer the quadrotor's lifted state from noisy position, velocity, attitude, and body-rate measurements. Although the sensor perturbations are modeled as Gaussian in their local coordinates, the nonlinear reconstruction can yield a non-Gaussian latent-state posterior. The Gaussian approximation may therefore become less accurate when attitude or translational-state uncertainty is large. The GECM guarantee concerns improvement of the conditional surrogate with the smoothing distribution held fixed; it does not establish convergence to a globally optimal model.

The finite-dimensional lifting and fixed-guide MPC prediction are designed for short-horizon quadrotor tracking. Prediction accuracy depends on the vehicle's angular motion and rollout horizon, while increasing the lifting degree incurs additional computational cost. Our control analysis establishes convexity and well-posedness of the stated optimization problems under their assumptions. The reported robustness is supported empirically under the evaluated disturbance and distribution-shift settings.

Finally, the experiments are conducted in MuJoCo Menagerie using noisy measurements of the full physical state and a common state estimator. Although the evaluation includes actuator dynamics, aerodynamic effects, external disturbances, and physical randomization, deployment on real vehicles requires validation with onboard sensing and computation. The Skydio X2 and Crazyflie experiments use separately trained models, so transfer without retraining remains untested. Physical-flight experiments and adaptation to changing operating conditions are natural directions for future work.