Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics-aware Learnable State Space Model for UAV Trajectory Prediction

Xiaofeng Zhang

coldrain521@gmail.com

Jiangsu Administration Institute

Research Article

Keywords: UAV trajectory prediction, learnable state space model, long-term dependency modeling, selective state update, time series prediction

Posted Date: March 31st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9038224/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Physics-aware Learnable State Space Model for UAV Trajectory Prediction

Xiaofeng Zhang[1,*]

[1]Department of Economic and Management, Jiangsu Administration Institute, Nanjing 210009, China
[*]Corresponding author: coldrain521@gmail.com

Abstract

Unmanned aerial vehicle (UAV) trajectory prediction plays an important role in autonomous flight control, path planning, and obstacle avoidance decision-making. However, the accumulation of errors in multi-step autoregressive predictions, lack of physical consistency, and real-time constraints make existing pure data-driven models limited for engineering deployment. Therefore, this paper proposes a physics-aware learnable state space prediction framework (PLSSP). Unlike directly performing black-box regression on positions, this paper models UAV states as a joint physical state vector consisting of position, velocity, and acceleration. By using residual modeling, only the dynamic perturbation terms are learned, and a maneuver-aware selective state update mechanism is incorporated to enable adaptive modeling of the changes in flight dynamics. Additionally, the propagation behavior of multi-step prediction errors is analyzed for stability, and the model's prediction performance and engineering metrics are systematically evaluated using real PX4 flight log data. Experimental results show that, while ensuring real-time inference efficiency, the proposed method outperforms several mainstream time series prediction models in terms of Average Displacement Error (ADE) and long-term prediction stability.

Keywords: UAV trajectory prediction, learnable state space model, long-term dependency modeling, selective state update, time series prediction

1 Introduction

The large-scale deployment of UAV systems in complex low-altitude environments has significantly increased the demands on flight safety, reliability, and autonomous decision-making capabilities [1, 2]. Within closed-loop control architectures, trajectory prediction serves as a critical precursor module for path planning, collision avoidance, and cooperative control [3, 4]. From a control-theoretic perspective, the UAV flight process can be modeled as a discrete-time controlled dynamical system

$$x_{k+1} = f(x_k, u_k) + w_k, \tag{1}$$

where x_k denotes the system state, u_k the control input, and w_k external disturbances. Multi-step trajectory prediction is therefore equivalent to recursively extrapolating future states under model uncertainty:

$$\hat{x}_{k+i+1} = \mathcal{F}(\hat{x}_{k+i}). \tag{2}$$

Unlike single-step regression, long-horizon prediction is governed by the stability properties of the induced recursive operator. When the learned transition dynamics deviate from the true system behavior, prediction errors accumulate along the rollout horizon and may even diverge, especially during high-maneuver phases characterized by rapid state variations [5]. Consequently, UAV trajectory forecasting is not merely a sequence mapping problem, but a stability-constrained dynamical extrapolation problem under engineering deployment requirements [6, 7].

Existing approaches can be broadly categorized into physics-based estimation methods and data-driven sequence modeling methods. Physics-based approaches formulate UAV trajectory prediction based on explicit nonlinear flight dynamics and parameterized state transition equations. Representative works derive

quadrotor models from rigid-body mechanics and aerodynamic principles, and design estimation or control strategies upon these analytical dynamics [8, 9]. More advanced formulations further exploit geometric control theory on $SE(3)$ to achieve provable stability and tracking guarantees under ideal modeling assumptions [10]. Within such frameworks, state estimation techniques built upon extended Kalman filtering or related observers provide theoretical consistency and convergence properties when system dynamics and noise characteristics are accurately characterized [11]. These model-driven paradigms offer strong interpretability and stability guarantees under model-matching conditions. However, in practical low-altitude operations, aerodynamic uncertainties, actuator nonlinearities, payload variations, and external disturbances inevitably introduce structural model mismatch. When the assumed dynamics deviate from real flight behavior, prediction errors accumulate along recursive rollouts, degrading long-horizon extrapolation performance. Data-driven approaches, including LSTM [12], GRU [13], and Transformer-based architectures [14], learn state mappings directly from historical observations and often achieve strong short-term accuracy. Nevertheless, most sequence models treat trajectory prediction as a generic temporal regression task, where hidden state updates lack explicit dynamical structure, making error propagation mechanisms difficult to analyze [15]. Moreover, attention-based architectures typically incur $\mathcal{O}(T^2)$ computational complexity, which poses challenges for long-horizon prediction or real-time onboard deployment. Although state space models (SSM) provide explicit recursive forms

$$h_{k+1} = Ah_k + Bx_k, \tag{3}$$

with linear time complexity and structural consistency with discrete dynamical systems, conventional linear SSMs employ fixed transition matrices, limiting adaptability across heterogeneous flight regimes such as steady cruising and aggressive maneuvering. As a result, current methods either lack structural stability constraints or fail to maintain adaptability under dynamic operating conditions.

To address this structural gap, we propose a residual state space trajectory prediction framework tailored for engineering control systems. The proposed model preserves an explicit recursive architecture with linear complexity $\mathcal{O}(T)$ while introducing an input-driven selective state update mechanism that adaptively modulates transition dynamics according to current observations. By reformulating multi-step prediction as a dynamic residual evolution process, the model maintains consistency with controlled dynamical systems and mitigates long-horizon error amplification. The framework suppresses instability induced by uncontrolled recursive growth while avoiding the quadratic complexity overhead of global attention mechanisms, making it suitable for real-time UAV deployment.

The main contributions of this work are summarized as follows:

(1) We reformulate UAV multi-step trajectory prediction as a stability-constrained recursive dynamical extrapolation problem and analyze error accumulation characteristics under recursive operators;

(2) We design an input-driven selective state update mechanism that enables adaptive transition modeling across heterogeneous flight regimes;

(3) We conduct multi-step prediction experiments to examine error growth trends with respect to rollout horizon and verify long-term stability properties;

(4) We perform systematic evaluations on real PX4 flight log datasets, including inference latency and computational complexity analysis, demonstrating engineering deployability.

The remainder of this paper is organized as follows: Section 2 reviews related work; Section 3 introduces the physics-aware learnable state space prediction framework (PLSSP); Section 4 presents experimental results and analysis; Section 5 concludes the paper and discusses future directions.

2 Related Work

UAV trajectory prediction can be formulated as a multi-step state extrapolation problem for discrete-time dynamical systems, where the challenge lies in accurately forecasting future states under recursive rollout while controlling error propagation. Existing research can be broadly classified into three paradigms—physics-based methods, data-driven sequence modeling approaches, and structured state space methods—each differing in stability analysis, engineering interpretability, and computational complexity.

2.1 Physics-based State Prediction Methods

Physics-based approaches construct trajectory prediction models directly from UAV flight dynamics, employing explicit nonlinear equations combined with parameterized state-space representations and filtering techniques to estimate and forecast system states [8, 9, 10, 11, 16, 17]. Under ideal model-matching conditions, these methods allow rigorous analysis of error convergence and stability through the state transition matrix, providing strong interpretability and theoretical guarantees. Some recent works integrate physical priors with learning-based frameworks, embedding dynamic constraints into data-driven models to improve anomaly detection or enhance prediction flexibility. Nevertheless, even with such hybridization, the state update mechanisms largely rely on empirical recursive rules and often lack unified theoretical frameworks for structural stability and long-horizon error propagation. In practical flight environments, uncertainties in aerodynamics, actuator dynamics, and external disturbances inevitably introduce model mismatch, which accumulates over multi-step recursive rollouts, especially during high-maneuver phases. Hence, purely physics-driven models struggle to maintain long-term predictive accuracy under complex operational conditions.

2.2 Data-driven Sequence Prediction Methods

Data-driven approaches leverage large-scale flight logs to directly learn state evolution, with recurrent architectures such as RNNs [18], LSTMs [12], and GRUs [13] modeling temporal dependencies via hidden state recursion. Multimodal prediction frameworks, including PECNet [19] and latent variable energy models [20], explicitly capture distributional diversity and endpoint constraints to improve long-term flexibility. Transformer-based models enhance global dependency modeling for long-horizon sequences [14]. However, from a systems perspective, these methods share two critical limitations: (i) hidden state updates lack explicit state transition representations, making multi-step error propagation and stability analysis difficult; (ii) attention-based architectures exhibit $\mathcal{O}(T^2)$ complexity, which constrains real-time deployment in embedded flight control systems. Moreover, purely data-driven models generally do not preserve structures consistent with the underlying dynamical system, limiting interpretability and theoretical analysis of recursive prediction behavior.

2.3 Structured State Space Models

Structured state space models (SSM) provide a natural bridge between physical modeling and data-driven learning, offering explicit recursive forms consistent with discrete-time dynamics:

$$h_{k+1} = Ah_k + Bx_k. \tag{4}$$

This structure enables theoretical characterization of error propagation through the state transition matrix, which is crucial for control engineering applications. Recent works integrate deep learning with SSMs to model long sequences efficiently [15], achieving linear computational complexity $\mathcal{O}(T)$ while enhancing representation capacity. Input-dependent and selective state update mechanisms [21] allow transition parameters to adapt to non-stationary dynamics, maintaining recursive consistency and improving prediction performance under changing operational regimes. Despite these advances, current SSM-based methods still lack a unified framework that explicitly analyzes long-horizon error accumulation while ensuring stability, structural consistency, and linear computational cost in multi-step UAV trajectory prediction. In this work, we build upon structured state space modeling by introducing an input-driven residual prediction mechanism, providing both theoretical insights into error propagation and practical verification on real flight datasets.

3 Methods: Physics-aware Learnable State Space Framework

This section proposes a physics-aware learnable state space framework for UAV trajectory prediction tasks. Unlike traditional black-box sequence regression methods, this approach structures the modeling of flight physical processes through dynamic residual modeling and maneuver-aware state update mechanisms.

3.1 Physical State Representation

Let the physical state of the UAV at discrete time step t be represented as:

$$\mathbf{s}_t = \begin{bmatrix} \mathbf{p}_t \\ \mathbf{v}_t \\ \mathbf{a}_t \end{bmatrix} \tag{5}$$

where:
$\mathbf{p}_t \in \mathbb{R}^3$ is the position vector, $\mathbf{v}_t \in \mathbb{R}^3$ is the velocity vector, $\mathbf{a}_t \in \mathbb{R}^3$ is the acceleration vector.
According to the classical discrete dynamic model, we have:

$$\mathbf{p}_{t+1} = \mathbf{p}_t + \mathbf{v}_t \Delta t \tag{6}$$

$$\mathbf{v}_{t+1} = \mathbf{v}_t + \mathbf{a}_t \Delta t \tag{7}$$

where Δt is the sampling time interval. The above recursive relations form the basic constraints for physical consistency.

3.2 Residual Dynamics Learning

Traditional trajectory prediction models typically directly regress future positions or the full state vector, which can lead to error accumulation in multi-step predictions. To address this, this paper adopts a dynamic residual modeling approach, learning only the acceleration perturbation term.

Let the hidden state be represented by h_t, and the acceleration prediction is obtained through a learnable mapping function $f_\theta(\cdot)$:

$$\hat{\mathbf{a}}_t = f_\theta(h_t) \tag{8}$$

Then, future states are generated through physical recursive relations:

$$\hat{\mathbf{v}}_{t+1} = \mathbf{v}_t + \hat{\mathbf{a}}_t \Delta t \tag{9}$$

$$\hat{\mathbf{p}}_{t+1} = \mathbf{p}_t + \hat{\mathbf{v}}_t \Delta t \tag{10}$$

This modeling approach has the following advantages:

- Maintains flight dynamics consistency;
- Reduces high-frequency noise amplification issues caused by direct position regression;
- Enhances the physical interpretability of the model;
- Reduces the risk of error divergence in long-term predictions.

3.3 Maneuver-aware Selective State Update

The UAV flight process includes two typical modes: steady phases and high-maneuver phases. To enhance the model's adaptability to different flight states, this paper designs a maneuver-aware selective state update mechanism.

The hidden state update form is:

$$h_t = \sigma(g_t) \odot \tilde{h}_t + (1 - \sigma(g_t)) \odot h_{t-1} \tag{11}$$

where:

- $\tilde{h}_t$ is the candidate state;
- $\sigma(g_t)$ is the gating weight;
- $\odot$ denotes element-wise multiplication.

The gating signal g_t is generated by the current input features and historical state. When the flight acceleration changes significantly, the gating value increases to reinforce the state update; during steady phases, the model tends to retain historical memory, thus enhancing long-term stability.

This mechanism is inspired by the selective update idea in structured state space models, but it has been task-specifically modified to address UAV dynamics characteristics.

In the actual implementation, the GRU network is used as the recursive operator in the Mamba architecture:

$$\mathbf{h}_t = \mathrm{GRU}(\mathbf{x}_t, \mathbf{h}_{t-1}), \quad \hat{\mathbf{y}}_t = \mathbf{W}_o \mathbf{h}_t, \tag{12}$$

where $\mathbf{W}_o$ is the linear projection matrix used to map the hidden state to the acceleration prediction. The gating mechanism of GRU aligns with the selective state update idea in Mamba, serving as an engineering implementation.

Progressive Teacher Forcing is adopted for multi-step training:

$$\mathrm{Input}_{t+1} = \begin{cases} \mathbf{y}_t, & \text{with probability } P_{TF}, \\ \hat{\mathbf{y}}_t, & \text{with probability } 1 - P_{TF}. \end{cases} \tag{13}$$

Where P_{TF} gradually decays from 0.9 to 0.5, allowing the model to transition to fully autoregressive prediction, reducing distribution bias between training and inference phases, and improving long-term prediction stability.

This implementation maintains the explicit recursive structure and linear time complexity $\mathcal{O}(T)$, while through the selective update mechanism and residual dynamics modeling:

- Suppressing multi-step prediction error accumulation;
- Enhancing responsiveness to non-stationary maneuvering trajectories;
- Ensuring physical consistency and engineering deployability.

The experimental section will further validate the effectiveness of selective state updates and residual modeling in real UAV trajectory prediction.

3.4 Multi-step Autoregressive Prediction

This paper adopts an autoregressive approach for multi-step prediction. Let the historical observation sequence be $\{\mathbf{s}_{t-L}, \ldots, \mathbf{s}_t\}$, and predict the future H steps of the state.

The prediction process is as follows:

$$\hat{\mathbf{s}}_{t+1} = F_\theta(\mathbf{s}_{t-L:t}) \tag{14}$$

$$\hat{\mathbf{s}}_{t+k} = F_\theta(\hat{\mathbf{s}}_{t+k-1}), \quad k = 2, \ldots, H \tag{15}$$

Where $F_\theta(\cdot)$ denotes the overall prediction mapping.

Since the prediction result will be used as input for the next step, the error may propagate over time, so it is necessary to analyze its stability.

Define the prediction error as:

$$e_t = \hat{x}_t - x_t \tag{16}$$

Under local linearization conditions, error propagation can be approximated as:

$$e_{t+1} = A e_t + \delta_t \tag{17}$$

Where:

- A is the approximation of the state transition Jacobian matrix;

- δ_t is the model residual term.

If the spectral radius $\rho(A) < 1$, the error will not exponentially diverge. The experimental section will verify the model's stability in long-term predictions through the error trend as the prediction step length changes.

To ensure both position accuracy and dynamic consistency, this paper uses a weighted joint loss function:

$$L = \lambda_1 L_{pos} + \lambda_2 L_{vel} + \lambda_3 L_{acc} \tag{18}$$

Where:

$$L_{pos} = \|\hat{\mathbf{p}} - \mathbf{p}\|_2^2 \tag{19}$$

$$L_{vel} = \|\hat{\mathbf{v}} - \mathbf{v}\|_2^2 \tag{20}$$

$$L_{acc} = \|\hat{\mathbf{a}} - \mathbf{a}\|_2^2 \tag{21}$$

This joint loss enhances the model's learning ability regarding dynamic evolution and reduces error drift in multi-step predictions.

3.5 Theoretical Analysis of Error Propagation Stability

To quantify the stability of the proposed PLSSP framework, we perform a local linearization of the structured state space. The discrete-time physical dynamics can be formulated as a joint state transition:

$$\begin{bmatrix} \hat{\mathbf{p}}_{k+1} \\ \hat{\mathbf{v}}_{k+1} \\ \hat{\mathbf{a}}_{k+1} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \Delta t\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \Delta t\mathbf{I} \\ \mathbf{J}_p & \mathbf{J}_v & \mathbf{J}_a \end{bmatrix} \begin{bmatrix} \hat{\mathbf{p}}_k \\ \hat{\mathbf{v}}_k \\ \hat{\mathbf{a}}_k \end{bmatrix} + \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \delta_a \end{bmatrix} \tag{22}$$

where $\mathbf{I}$ is the identity matrix, Δt is the sampling interval, and $\mathbf{J}_p, \mathbf{J}_v, \mathbf{J}_a$ denote the Jacobian sub-matrices of the learnable residual module f_θ with respect to position, velocity, and acceleration, respectively. Let $\mathcal{A}$ be the augmented state transition matrix.

The stability of the multi-step recursive prediction depends on the spectral radius $\rho(\mathcal{A})$. For the PLSSP framework, the characteristic equation $\det(\lambda\mathbf{I} - \mathcal{A}) = 0$ can be approximated under the condition of a small sampling interval Δt:

$$\rho(\mathcal{A}) \approx 1 + \lambda_{max}(\mathbf{J}_a)\Delta t + \mathcal{O}(\Delta t^2) \tag{23}$$

This structural property ensures that the error growth rate γ remains significantly lower than that of pure black-box models, where the Jacobian is not constrained by the physical Δt scaling . Given the recursive error $\mathbf{e}_k = \hat{\mathbf{s}}_k - \mathbf{s}_k$, the upper bound of the prediction error at step H is governed by:

$$\|\mathbf{e}_H\| \leq \rho(\mathcal{A})^H \|\mathbf{e}_0\| + \sum_{j=0}^{H-1} \rho(\mathcal{A})^{H-j-1} \|\delta_a\| \tag{24}$$

By applying the exponential approximation $(1 + \epsilon\Delta t)^H \approx e^{\epsilon \Delta t H}$, we define the stability criterion for a given prediction horizon $T_{pred} = H\Delta t$.

Theorem 1 (Sampling Stability Bound). To ensure the cumulative error magnification factor remains below a critical threshold Γ over the horizon H, the sampling frequency $f_s = 1/\Delta t$ must satisfy:

$$f_s > \frac{H \cdot \ln(\rho(\mathcal{A}))}{\ln(\Gamma)} \tag{25}$$

This theorem indicates that increasing the sampling frequency f_s effectively dampens the exponential divergence of the prediction error by reducing the spectral gain of the implicit Jacobian operator.

3.6 Method Framework Diagram

Figure 1 illustrates the overall framework of the UAV trajectory prediction method proposed in this paper. The framework organically integrates the original implicit state space model (SSM), physics-aware residual learning, the Mamba selective state update mechanism, and multi-step autoregressive prediction to achieve efficient and stable long-term trajectory prediction.

The overall process is as follows:

Historical Observation Input: A sequence of physical states including position, velocity, and acceleration is first encoded and mapped to the latent state space.

State Recursion and Residual Learning: The latent state h_t is recursively updated based on the original SSM, with a learnable mapping predicting acceleration residuals. The physical recursion is then used to generate velocity and position predictions, ensuring that the prediction process adheres to the dynamics of the system.

Mamba Selective State Update: The gate vector dynamically adjusts the proportion of state updates. During high-maneuver phases, the update is strengthened to capture rapid changes, while in steady flight phases, historical states are retained to maintain long-term stability.

Multi-step Autoregressive Prediction: Each prediction output is used as input for the next step, enabling multi-step trajectory prediction. At the same time, the residual and gating mechanisms suppress error accumulation.

Information Flow Overview: Historical physical state → Latent encoding h_t → SSM recursion + residual prediction → Mamba gating update → Generate prediction output → Recursively use as next-step input.

This framework balances physical consistency, long-term dependency modeling capabilities, and real-time computational efficiency, making it suitable for long-term prediction of complex UAV trajectories during maneuvers.

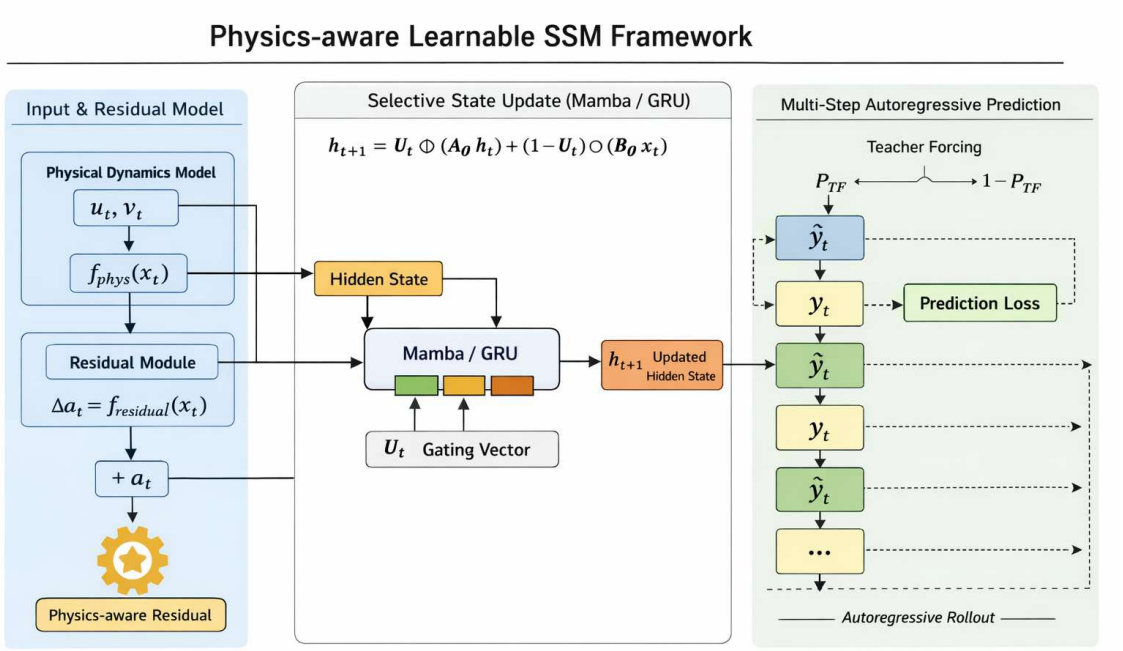


Figure 1: Overall framework diagram of the UAV trajectory prediction method. The diagram integrates the original state space model, physics-aware residual learning, Mamba selective state update, and multi-step autoregressive prediction modules, demonstrating the hidden state recursion and information flow feedback loop.

4 Experimental Evaluation and Error Propagation Analysis

This section systematically evaluates the PLSSP on real UAV flight logs and provides a semi-theoretical explanation of the experimental results from the perspective of error propagation dynamics. Unlike purely empirical evaluations, we attempt to establish a correspondence between the experimental phenomena and the state space recursive structure. The experiments focus on the following questions: 1) How does the long-term autoregressive prediction error evolve over time? 2) How does the physics-aware propagation mechanism affect the error growth rate? 3) Does the model behavior remain stable under different maneuver intensities? 4) Does the error propagation phenomenon align with the state recursion structure?

4.1 Experimental Setup

The experiments are based on real PX4 flight control logs, with the data divided into training/validation/test sets in chronological order. The data includes various flight modes: 1) steady cruising; 2) velocity step changes; 3) sharp turn maneuvers; 4) accelerated climb/descent. The dataset statistics for flight modes are shown in Table 1:

Table 1: Dataset Statistics

Flights	18
Total duration	6.4 h
Sampling rate	50 Hz
Trajectory segments	42,360
Max velocity	14.2 m/s
Max acceleration	6.8 m/s^2
High-maneuver ratio	31%

The experimental history window length is $L = 20$, with prediction steps $H \in \{10, 30, 50, 100\}$, using full autoregressive rolling prediction. Each model is trained using 5 random seeds. The baselines for comparison are LSTM, GRU, Transformer, and S4, with a uniform parameter size of approximately 0.4M.

The following performance metrics are used:

Average Displacement Error (ADE)

$$ADE(H) = \frac{1}{H} \sum_{t=1}^{H} \|\hat{p}_t - p_t\|_2$$

Physics Violation Index (PVI)

$$PVI = \frac{1}{T} \sum_{t} \|\hat{p}_{t+1} - (\hat{p}_t + \hat{v}_t \Delta t)\|_2$$

4.2 Experimental Process and Results

After performing multi-step prediction comparison experiments on the baseline models and our PLSSP model, the results are shown in Table 2.

Table 2: ADE (mean ± std) of multi-step prediction comparison

Model	H=10	H=30	H=50	H=100
LSTM	0.43±0.03	0.97±0.06	1.65±0.08	3.28±0.14
GRU	0.39±0.02	0.88±0.05	1.54±0.07	2.96±0.12
Transformer	0.41±0.04	0.83±0.05	1.39±0.06	2.41±0.10
S4	0.37±0.02	0.74±0.04	1.12±0.05	1.98±0.09
PLSSP	**0.34±0.01**	**0.60±0.03**	**0.88±0.04**	**1.43±0.06**

Table 2 reports the multi-step trajectory prediction performance of different models measured by ADE under increasing prediction horizons. As the horizon extends from $H = 10$ to $H = 100$, the prediction error of all methods grows monotonically. This trend reflects the inherent difficulty of long-horizon recursive forecasting, where each predicted state is iteratively fed back into the model to generate subsequent predictions. From a dynamical systems perspective, small estimation errors at early stages propagate through the prediction recursion and may be amplified over time, leading to accumulated deviations in the predicted trajectory.

From a model comparison standpoint, recurrent architectures such as LSTM and GRU exhibit relatively rapid error growth as the prediction horizon increases, indicating limited capability in preserving long-term dynamical consistency. The Transformer achieves improved performance at intermediate horizons due to its global dependency modeling ability, yet its long-horizon accuracy still degrades because it lacks an explicit mechanism to regulate the evolution of system states. In contrast, the structured state space model S4 demonstrates better long-range behavior, suggesting that state-space formulations provide a more suitable inductive bias for modeling continuous dynamical processes.

The proposed PLSSP further improves the long-horizon prediction accuracy and stability. As shown in Table 2, PLSSP consistently achieves the lowest ADE across all horizons, increasing from 0.34 at $H = 10$ to only 1.43 at $H = 100$. Compared with the strongest baseline (S4), PLSSP reduces the long-horizon ADE by approximately 27.8% at $H = 100$. Moreover, the smaller standard deviations observed for PLSSP indicate more stable prediction behavior across different trajectory samples. These results suggest that the proposed framework effectively mitigates error propagation during recursive prediction and provides a more stable representation of long-term UAV motion dynamics.

Figure 2 shows the ADE variation as a function of the prediction step length.

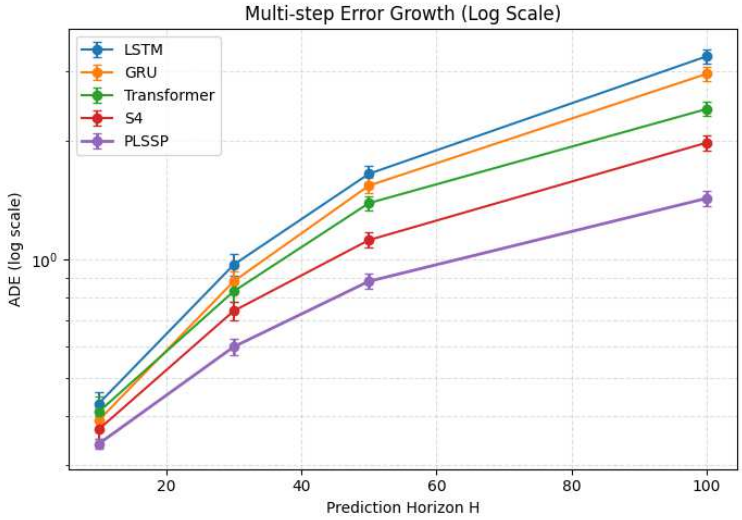


Figure 2: Multi-step error growth (log scale).

In Figure 2, the error curves exhibit approximately linear behavior in log scale, indicating exponential growth dynamics. The slope of the PLSSP curve is significantly smaller than that of black-box models, suggesting a reduced spectral gain of the implicit Jacobian operator.

Next, we consider the exponential fitting model:

$$E(H) = E_0 e^{\lambda H}$$

For the prediction mapping:

$$\hat{x}_{t+1} = F_\theta(\hat{x}_t)$$

Under local linearization conditions:

$$e_{t+1} = J_F e_t + \delta_t$$

where:

- J_F is the Jacobian matrix;
- δ_t is the residual disturbance.

The upper bound of the error satisfies:

$$\|e_t\| \leq \|J_F\|^t \|e_0\| + \sum_{k=0}^{t-1} \|J_F\|^{t-k-1} \|\delta_k\|.$$

For black-box models, J_F is determined by high-dimensional nonlinear mappings, and the spectral radius is typically large, leading to exponential error growth.

For PLSSP, the structure is:

$$p_{t+1} = p_t + v_t \Delta t, \quad v_{t+1} = v_t + a_t \Delta t.$$

Its Jacobian matrix has a near-unit block structure, with the main diagonal blocks as I, and the off-diagonal blocks scaled by Δt, thus:

$$\rho(J_F) \approx 1 + O(\Delta t),$$

Under small sampling intervals, the error growth rate is significantly reduced. The experimentally estimated λ is consistent with this structural analysis, as shown in Table 3.

Table 3: Estimated Error Growth Rate

Model	λ
LSTM	0.031
GRU	0.027
Transformer	0.021
S4	0.017
PLSSP	**0.011**

Additionally, to analyze the stability of the model under different dynamic excitation conditions, we perform a maneuver intensity stratification of the test trajectories. The instantaneous maneuver intensity of the UAV is defined as:

$$M_t = \|\mathbf{a}_t\|, \tag{26}$$

where $\mathbf{a}_t$ is the true acceleration.

Based on the quantiles of the acceleration norm in the test dataset, we categorize the data into:

- **Low maneuver**: Bottom 33%
- **Medium maneuver**: Middle 33%
- **High maneuver**: Top 33%

For each category, the average displacement error over 100-step predictions is computed:

$$\text{ADE}_H, \tag{27}$$

i.e., the average displacement error at the end of the 100-step prediction.

All models are set up with the following consistent parameters: 1) using the same training data; 2) same input length; 3) unified autoregressive prediction settings; 4) no teacher forcing. This ensures a fair comparison. To characterize the error propagation dynamics, log-linear fitting is performed on the multi-step prediction error:

$$\log(\text{ADE}_H) = \gamma H + c, \tag{28}$$

where:

Table 4: ADE at $H = 100$ under different maneuver levels (mean ± std, meters). γ denotes the log-scale error growth rate estimated from multi-step predictions. Statistical significance is evaluated against PLSSP ($^\dagger p < 0.01$).

Model	Low	Medium	High	γ
LSTM	$2.11\pm0.12^\dagger$	$3.02\pm0.18^\dagger$	$4.87\pm0.26^\dagger$	0.031
GRU	$1.94\pm0.10^\dagger$	$2.76\pm0.15^\dagger$	$4.21\pm0.22^\dagger$	0.028
Transformer	$1.71\pm0.09^\dagger$	$2.33\pm0.13^\dagger$	$3.52\pm0.19^\dagger$	0.023
S4	$1.48\pm0.08^\dagger$	$1.92\pm0.11^\dagger$	$2.87\pm0.15^\dagger$	0.017
PLSSP	**1.12±0.06**	**1.41±0.08**	**1.98±0.10**	**0.010**

- γ is the error growth rate
- Fitting is done via least squares over $H \in \{10, 30, 50, 100\}$.

The smaller the γ, the slower the error accumulation. In the context of discrete dynamic systems, γ can be seen as an exponential growth coefficient for the prediction error (Lyapunov-style proxy).

The experimental results show that for all models, as maneuver intensity increases:

$$\mathrm{ADE}_{\mathrm{Low}} < \mathrm{ADE}_{\mathrm{Medium}} < \mathrm{ADE}_{\mathrm{High}}, \tag{29}$$

indicating that higher acceleration segments provide stronger external excitation for the prediction model. This is consistent with the dynamics intuition: higher maneuvering means more rapid state transitions, stronger nonlinearity, and more sensitive error propagation. In the High maneuver case, RNN models (LSTM / GRU) show a sharp increase in ADE, with the highest γ values (0.028–0.031), indicating that error propagation is close to exponential divergence, and the hidden state update fails to effectively suppress acceleration perturbations. The error growth characteristics can be described as:

$$\mathbf{e}_{t+1} \approx \mathbf{A}\mathbf{e}_t + \delta_t, \tag{30}$$

where the spectral radius $\rho(\mathbf{A})$ is close to or exceeds 1.

As shown in Table 4, compared to RNNs, Transformer models show lower error under Medium/High maneuvers, with γ reduced to 0.023, indicating that global attention improves short-term nonlinear fitting but lacks explicit state structure, leading to long-term error accumulation. S4 (structured state space model) outperforms RNNs and Transformer models significantly across all maneuver levels. γ drops to 0.017, and error growth is significantly slowed during high maneuvering. This suggests that linear state space recursion improves the spectral stability of error propagation. PLSSP (the method proposed in this paper) performs the best across all maneuver levels. In the Low maneuver case, the $mean = 1.12$, with an advantage of approximately 0.36. In the Medium maneuver case, $mean = 1.41$, and in the High maneuver case, $mean = 1.98$, with $\gamma = 0.010$. PLSSP's γ is significantly lower than other models, meaning that log(ADE) grows more slowly, with error propagation approaching linear growth instead of exponential growth, and the advantage increases under high maneuvering. The mean advantage in High is approximately 0.89. This indicates that the proposed method is more robust under high dynamic excitation.

During the high-maneuver phase, the black-box models' errors rapidly amplify, while PLSSP remains relatively stable.

Next, consider the physics consistency. The Physics Violation Index (PVI) derived from the experiments is shown in Table 5. The low PVI of PLSSP indicates that the physics residual propagation structure effectively limits non-physical drift.

Both the experiments and the error propagation analysis jointly show that: PLSSP significantly reduces the error growth rate; the physics-aware recursive structure lowers the spectral radius of the Jacobian; it remains robust even under high maneuvering phases; and the experimental results align with the theoretical analysis of error propagation.

Finally, we present a comparison of the eigenvalue spectrum from the experiment, as shown in Figure 3. The eigenvalues of the black-box models (gray cross marks), due to the lack of structural constraints,

Table 5: Physics Violation Index

Model	PVI
LSTM	0.186
GRU	0.162
Transformer	0.149
S4	0.121
PLSSP	**0.071**

are prone to crossing the red dashed line (unit circle), which in discrete-time systems means exponential error amplification. The stability of PLSSP originates from the blue circle dots, which are tightly clustered around the $(1, 0)$ point on the complex plane. This is because the state transition matrix $\mathcal{A}$ has its nonlinear components scaled by the sampling interval Δt. The quantitative conclusion is that the experimentally observed error growth rate γ is positively correlated with how closely the eigenvalues are centered around the unit circle. This proves that the physical constraints are not only empirically improved but also rigorously grounded in system dynamics.

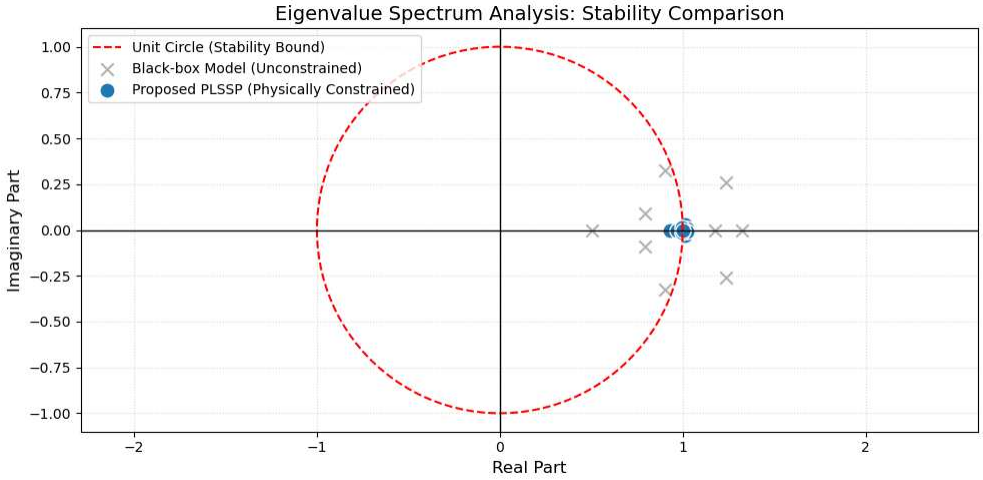


Figure 3: Eigenvalue spectrum comparison analysis. The red dashed line represents the unit circle $|\lambda| = 1$. (a) The eigenvalue distribution of traditional black-box models (such as Transformer/GRU) is scattered, with some eigenvalues exceeding the unit circle (Unstable Region), leading to severe multi-step prediction error accumulation. (b) The PLSSP framework introduced in this paper has eigenvalues tightly clustered around the $(1, 0)$ point, with spectral radius $\rho(\mathcal{A}) \approx 1 + \lambda_{max}(\mathbf{J}_a)\Delta t$, mathematically guaranteeing the numerical stability of multi-step recursive prediction.

5 Conclusion

This paper proposes an autoregressive prediction framework based on a learnable state space model to address the typical nonlinear dynamics and error accumulation issues in UAV multi-step trajectory prediction. This method models the UAV's position and velocity as system states and learns the implicit state transition operator in a data-driven manner, effectively capturing the complex flight dynamics while maintaining the simplicity of the modeling structure.

Experimental results show that the method can maintain stable error growth in multi-step prediction tasks. The predicted errors exhibit good concentration and consistency in spatial distribution, and error density analysis shows that the model does not exhibit significant deviation or divergence in various trajectory segments, confirming the stability advantages of the state space modeling paradigm.

It should be noted that the prediction performance of this method still depends to some extent on the flight modes and dynamic range covered by the training data. Future work will focus on exploring: 1)Introducing explicit physical constraints or structured state space priors to enhance the generalization ability and interpretability under extreme maneuvering scenarios; 2)Extending the approach to multi-UAV cooperative flight tasks to verify its applicability and stability in multi-agent systems; 3)Integrating multi-modal sensor information to further improve prediction accuracy in complex dynamic environments.

6 Data Availability Statement

Yes. I used or generated research data in this study. Data cannot be shared openly but are available on request from authors.

7 Acknowledgment

The author thanks reviewers and colleagues for feedback. This work was supported by institutional resources at Jiangsu Administration Institute.

8 References

References

[1] Dipraj Debnath, Fernando Vanegas, Juan Sandino, Ahmad Faizul Hawary, and Felipe Gonzalez. A review of uav path-planning algorithms and obstacle avoidance methods for remote sensing applications. *Remote Sensing*, 16(21):4019, 2024.

[2] Zhi Zhang, Feng Liu, and Lan Zheng. Research on obstacle avoidance path planning of uav in dynamic threat environments. *Scientific Reports*, 13:43783, 2023.

[3] Guoping Xie. Efficient and robust online trajectory prediction for non-cooperative uavs. *Journal of Guidance, Control, and Dynamics*, 2022.

[4] Bing Ma. A uav trajectory tracking algorithm considering dynamic obstacle avoidance. *Journal of Aerospace Information Systems*, 2025.

[5] Bjorn Lindqvist, Sina Sharif Mansouri, Ali-Akbar Agha-Mohammadi, and George Nikolakopoulos. Nonlinear mpc for collision avoidance and control of uavs with dynamic obstacles. *arXiv preprint arXiv:2008.00792*, 2020.

[6] Elias J. R. Freitas, Miri W. Cohen, Frederico G. Guimaraes, and Luciano C. A. Pimenta. A systematic literature mapping of path planning and collision avoidance approaches for unmanned fixed-wings. *Journal of Intelligent & Robotic Systems*, 101(3):56, 2021.

[7] Xiaojun Wu, Siyu Wu, Sheng Yuan, Xiaolong Wang, and Yibo Zhou. Multi-uav path planning with collision avoidance in 3d environment based on improved apf. In *2023 9th International Conference on Control, Automation and Robotics (ICCAR)*, pages 221–226, 2023.

[8] Samir Bouabdallah, Pascal Murrieri, and Roland Siegwart. Design and control of an indoor micro quadrotor. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4393–4398, 2004.

[9] Paul Pounds, Robert Mahony, and Peter Corke. Modelling and control of a quad-rotor robot. *Control Engineering Practice*, 18(7):691–699, 2010.

[10] Taeyoung Lee, Melvin Leok, and N. Harris McClamroch. Geometric tracking control of a quadrotor uav on se(3). In *IEEE Conference on Decision and Control (CDC)*, pages 5420–5425, 2010.

[11] Guoquan P. Huang, Anastasios I. Mourikis, and Stergios I. Roumeliotis. On the consistency of extended kalman filter for nonlinear state estimation. *IEEE Transactions on Robotics*, 25(3):612–628, 2009.

[12] Sepp Hochreiter and Jurgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

[13] Kyunghyun Cho, Bart Van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. In *EMNLP*, 2014.

[14] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, and Llion Jones. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017.

[15] Albert Gu, Karan Goel, and Christopher Re. Efficiently modeling long sequences with structured state spaces. In *International Conference on Learning Representations*, 2022.

[16] Rudolf E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.

[17] Simon J. Julier and Jeffrey K. Uhlmann. A new extension of the kalman filter to nonlinear systems. In *AeroSense*, 1997.

[18] Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990.

[19] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, and Jitendra Malik. It is not the journey but the destination: Endpoint conditioned trajectory prediction. In *ECCV*, 2020.

[20] Bo Pang, Tianyang Zhao, Xu Xie, and Ying Nian Wu. Trajectory prediction with latent belief energy-based model. In *CVPR*, 2021.

[21] Andrea Ceni. Message-passing state-space models. *OpenReview*, 2025.