FROM HISTORICAL ERRORS TO FUTURE TRAJECTORIES: LEARNING TRANSFERABLE KNOWLEDGE FOR SATELLITE ORBIT PREDICTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Satellite orbit prediction is fundamental to modern space operations, with an increasing demand for accurate trajectory forecasts driven by the rapid expansion of satellite constellations. Recent studies employ deep learning to model the residuals between predicted and true trajectories, providing a promising solution for orbit forecasting. However, such residual learning paradigms simply treat prediction errors as supervision, overlooking the inherent periodicity of orbital motion and thereby limiting their predictive performance. To this end, we introduce ORBITECHO, a novel residual correction framework that encodes historical prediction experience into *transferable knowledge* for accurate trajectory forecasting. Specifically, ORBITECHO first stores forecasting errors from preceding orbital cycles in a memory bank. It then retrieves phase-aligned trajectory errors at similar orbital states as informative references for residual correction. For each retrieved candidate, ORBITECHO further exploits orbital periodicity to model its residual evolution over one orbital cycle, which enables historical errors to be adapted to the current prediction. Together, ORBITECHO distills historical predictions into two complementary forms of transferable knowledge, *i.e.*, *trajectory errors* and *residual dynamics*, which are jointly leveraged to refine orbit forecasts. Extensive experiments on real-world satellites demonstrate the superior performance of ORBITECHO, reducing the average MAE from 8.40 km for a state-of-the-art learning-based baseline to 0.31 km, yielding an order-of-magnitude improvement. ORBITECHO also exhibits strong generalization to unseen satellites and robust performance under severe geomagnetic storms.

1 INTRODUCTION

Satellite orbit prediction aims to forecast future trajectories from historical observations and is essential for modern space operations, including mission planning (Zhang et al., 2024; Fan et al., 2025) and collision avoidance (Reiland et al., 2021; Caldas & Soares, 2024b). With the rapid expansion of satellite constellations and growing congestion in orbital environment (European Space Agency, 2025), the demand for accurate trajectory forecasts continues to increase. Uncertainties in orbital states and imperfect modeling of perturbing forces, such as the Earth's nonspherical gravity field and atmospheric drag, can cause the accumulation of prediction errors and result in kilometer-scale deviations (Caldas & Soares, 2024a; Licata & Mehta, 2022; Caldas & Soares, 2024b; Fan et al., 2024). Traditional physics-based propagators, *e.g.*, Simplified General Perturbations 4 (SGP4) (Hoots & Roehrich, 1980), predict future orbital states based on predefined dynamical models, whose accuracy heavily depends on the quality of initial states and perturbation modeling.

In recent years, learning-based approaches have emerged as a promising alternative, which leverages large-scale orbital observations to learn orbital evolution directly from data. Driven by advances in time-series forecasting (Zeng et al., 2023; Wu et al., 2023; Liu et al., 2024; Chen et al., 2026), diverse architectures (Lin & Chen, 2025; Lin et al., 2025b; Shin et al., 2022; Jeong et al., 2024) have been devised to capture complex temporal dependencies in orbital trajectories while reducing reliance on explicit modeling of perturbing forces. Beyond direct trajectory forecasting, a growing line of research has shifted toward a two-stage residual correction paradigm (Lin et al., 2025a; 2026; Liao

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.

et al., 2026; Zhang et al., 2026). Under this framework, a base predictor first generates a reference trajectory, which is then refined by a learning-based model that estimates its prediction errors.

Despite this progress, existing residual correction solutions typically consider prediction errors as supervision, while neglecting the inherent periodicity of orbital motion. Consequently, the patterns embedded in historical errors are not fully exploited, thereby leading to unsatisfactory forecasting results. This motivates us to delve deeper into historical prediction errors and explore how past prediction experience can be leveraged to facilitate accurate trajectory correction.

To this end, we propose ORBITECHO, a residual correction framework tailored for satellite orbit prediction. It encodes historical prediction experience into *transferable knowledge* for accurate trajectory forecasting. Specifically, we first construct a trajectory memory bank that records historical orbital states and their corresponding forecasting errors. Leveraging orbital periodicity, phase-aligned residual matching is then performed to retrieve historical cases with similar orbital states, whose errors provide informative references for trajectory correction. To model such evolution, ORBITECHO further traces each retrieved candidate across adjacent orbital cycles and captures the corresponding residual changes. The resulting cross-cycle error evolution characterizes the residual dynamics, which enables the retrieved errors to be adapted to the current prediction. Together, the retrieved trajectory errors and learned residual dynamics constitute *transferable knowledge*. Their transferability is further supported by theoretical analysis of phase-aligned prediction errors and cross-cycle residual dynamics. By capturing both what errors occurred under similar orbital conditions and how these errors evolve across orbital cycles, ORBITECHO effectively leverages historical prediction experience for accurate trajectory correction.

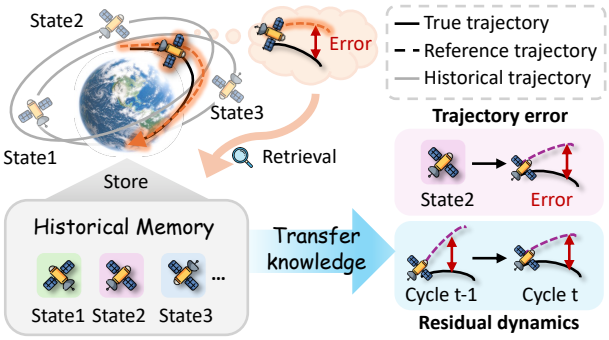


Figure 1: Illustration of transferable knowledge from historical prediction experience. Historical trajectory errors capture recurring error patterns under similar orbital conditions, while residual dynamics characterize their evolution across orbital cycles.

Extensive experiments on real-world orbital traces from three in-orbit satellites, *i.e.*, Sentinel-3A, Sentinel-3B, and SARAL, demonstrate that ORBITECHO surpasses both general time-series forecasting models and specialized orbit prediction methods. Notably, ORBITECHO achieves an average position error of only 311.4 m at a 60-minute prediction horizon, representing an order-of-magnitude improvement over learning-based baselines. Beyond in-domain evaluation, ORBITECHO also exhibits strong cross-satellite generalization to two unseen satellites. Moreover, experiments under severe geomagnetic storm conditions further highlight its robustness to challenging space weather.

In a nutshell, our contributions are three-fold:

- We propose ORBITECHO, a novel residual correction framework that encodes historical prediction experience into transferable knowledge for accurate trajectory forecasting.
- We develop phase-aligned residual matching and residual dynamics modeling to capture what errors occurred under similar orbital conditions and how they evolve across orbital cycles, with theoretical analysis supporting their transferability.
- Extensive experiments on real-world satellites demonstrate that ORBITECHO achieves superior forecasting accuracy and generalizes well to unseen satellites without fine-tuning.

2 METHODOLOGY

2.1 OVERVIEW OF ORBITECHO

Given a historical observation sequence $\mathbf{X}_t = [\mathbf{x}_{t-L+1}, \ldots, \mathbf{x}_t] \in \mathbb{R}^{L \times D}$ of length L, satellite orbit prediction aims to forecast the future trajectory $\mathbf{Y}_t = [\mathbf{x}_{t+1}, \ldots, \mathbf{x}_{t+H}] \in \mathbb{R}^{H \times D}$ over a prediction

horizon of length H. Each orbital state is represented as $\mathbf{x}_t = [\mathbf{p}_t, \mathbf{v}_t]$, where $\mathbf{p}_t = (x_t, y_t, z_t)$ and $\mathbf{v}_t = (v_t^x, v_t^y, v_t^z)$ denote the satellite position and velocity, with $D = 6$. Following the prediction-then-correction" paradigm, a base predictor f_θ first generates a reference trajectory:

$$\mathbf{Y}_t^{\text{ref}} = f_\theta(\mathbf{X}_t) = [\mathbf{x}_{t+1}^{\text{ref}}, \ldots, \mathbf{x}_{t+H}^{\text{ref}}]. \tag{1}$$

The resulting prediction residual is represented in the radial-transverse-normal (RTN) frame as $\mathbf{R}_t = [\mathbf{r}_{t+1}, \ldots, \mathbf{r}_{t+H}] \in \mathbb{R}^{H \times D}$, where $\mathbf{r}_{t+h}$ denotes the residual between the true state $\mathbf{x}_{t+h}$ and its reference prediction $\mathbf{x}_{t+h}^{\text{ref}}$ at the h-th prediction step.

As illustrated in Figure 2, building upon the reference trajectory, ORBITECHO first performs phase-aligned residual matching via a trajectory memory bank to retrieve cases with similar orbital states, and then traces the retrieved candidates across adjacent orbital cycles to model their cross-cycle residual changes. A router aggregates these residual dynamics and combines them with a phase-aligned error anchor to produce the residual correction $\hat{\mathbf{R}}t = [\hat{\mathbf{r}}t + 1, \ldots, \hat{\mathbf{r}}_{t+H}]$. The correction is finally transformed to the Cartesian frame to obtain the corrected trajectory $\hat{\mathbf{Y}}_t$.

2.2 RETRIEVING TRANSFERABLE TRAJECTORY ERRORS

Historical prediction errors provide informative correction cues when similar orbital states recur, whereas arbitrary historical errors may be unreliable. We thus construct a trajectory memory bank:

$$\mathcal{M}_t = \left\{ \left(\mathbf{Y}_i^{\text{ref}}, \mathbf{R}_i \right) \,\middle|\, t - 3P \le i < t \right\}, \tag{2}$$

where P is the orbital period estimated from the reference trajectory, $\mathbf{Y}_i^{\text{ref}}$ and $\mathbf{R}_i$ denote the reference trajectory and residual sequence of historical case i, respectively. In our study, we retain the most recent three orbital cycles to preserve recurring patterns while limiting the influence of outdated prediction errors from distant cycles.

To identify relevant historical errors, we measure the phase similarity between the current reference state and historical reference states. For the h-th prediction step, the orbital phase of $\mathbf{x}_{t+h}^{\text{ref}}$ is represented using the normalized directions of its position and velocity:

$$\boldsymbol{\phi}_{t+h} = \left[\frac{\mathbf{p}_{t+h}^{\text{ref}}}{\|\mathbf{p}_{t+h}^{\text{ref}}\|_2}, \frac{\mathbf{v}_{t+h}^{\text{ref}}}{\|\mathbf{v}_{t+h}^{\text{ref}}\|_2} \right], \tag{3}$$

where $\mathbf{p}_{t+h}^{\text{ref}}$ and $\mathbf{v}_{t+h}^{\text{ref}}$ denote the position and velocity of $\mathbf{x}_{t+h}^{\text{ref}}$, respectively. Similarly, $\boldsymbol{\phi}_{i+h}$ is the phase representation of $\mathbf{x}_{i+h}^{\text{ref}}$ from historical case i. Their phase similarity is measured as:

$$c_{t,i}^{(h)} = \frac{1}{2} \boldsymbol{\phi}_{t+h}^\top \boldsymbol{\phi}_{i+h}, \tag{4}$$

For each prediction step, we then retrieve the K historical cases with the highest phase similarities from the trajectory memory bank:

$$\mathcal{N}_{t,h} = \text{TopK}_{i \in [t-3P, t-h]} \left(c_{t,i}^{(h)} \right), \tag{5}$$

where $\mathcal{N}_{t,h}$ is the set of retrieved historical cases for the h-th prediction step. The constraint $i+h \le t$ ensures that the corresponding historical residual is available at current prediction time.

2.3 LEARNING TRANSFERABLE RESIDUAL DYNAMICS

Although phase-aligned trajectory errors provide informative correction references, their direct reuse assumes that prediction errors remain stable across orbital cycles. To capture their evolution, we model the residual variation between successive cycles as residual dynamics.

For each retrieved candidate $i \in \mathcal{N}_{t,h}$, we identify its phase-aligned predecessor in the preceding orbital cycle. Since the orbital period may vary over time, the predecessor is selected from a range of candidate cycle lengths $\mathcal{P}$:

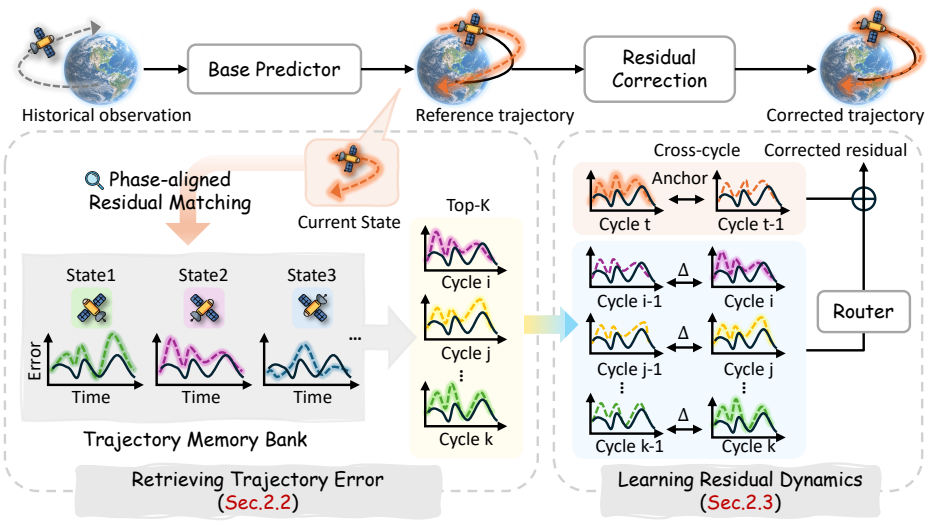


Figure 2: Overview of ORBITECHO. Given a reference trajectory from a base predictor, ORBITECHO retrieves phase-aligned historical cases and models their cross-cycle residual dynamics. A router further aggregates these dynamics and combines them with a phase-aligned error anchor to produce the final residual correction.

$$\pi(i,h) = i - \arg\max_{p\in\mathcal{P}} \frac{1}{2}\boldsymbol{\phi}_{i+h}^{\top}\boldsymbol{\phi}_{i-p+h}. \tag{6}$$

The corresponding cross-cycle residual change is given by:

$$\Delta\mathbf{r}_{i+h} = \mathbf{r}_{i+h} - \mathbf{r}_{\pi(i,h)+h}. \tag{7}$$

where $\Delta\mathbf{r}_{i+h}$ describes the evolution of prediction errors across adjacent orbital cycles at similar phases. As retrieved candidates may exhibit distinct residual dynamics, we introduce a query-key router to adaptively determine their contributions. Each candidate is represented by a key $\mathbf{k}_{i,h}$ derived from its residual dynamics and orbital information, while the query $\mathbf{q}_{t,h}$ is constructed from the current reference state and prediction horizon. The routing weights are defined as:

$$\alpha_{t,h,i} = \frac{\exp\left(\mathbf{q}_{t,h}^{\top}\mathbf{k}_{i,h}/\sqrt{d}\right)}{\sum_{j\in\mathcal{N}_{t,h}}\exp\left(\mathbf{q}_{t,h}^{\top}\mathbf{k}_{j,h}/\sqrt{d}\right)}, \qquad i\in\mathcal{N}_{t,h}, \tag{8}$$

where d is the dimension of the query and key representations. The historical residual dynamics are finally aggregated according to their routing weights:

$$\widehat{\Delta\mathbf{r}}_{t+h} = \sum_{i\in\mathcal{N}_{t,h}} \alpha_{t,h,i}\Delta\mathbf{r}_{i+h}. \tag{9}$$

We similarly identify a phase-aligned predecessor for the current prediction:

$$\pi(t,h) = t - \arg\max_{p\in\mathcal{P}} \frac{1}{2}\boldsymbol{\phi}_{t+h}^{\top}\boldsymbol{\phi}_{t-p+h}. \tag{10}$$

Its residual $\mathbf{r}_{\pi(t,h)+h}$ serves as an error anchor from the preceding orbital cycle. Combining this anchor with the transferred residual dynamics yields the final residual correction:

$$\hat{\mathbf{r}}_{t+h} = \mathbf{r}_{\pi(t,h)+h} + \widehat{\Delta\mathbf{r}}_{t+h}. \tag{11}$$

Applying this procedure across all prediction steps gives $\hat{\mathbf{R}}_t$, which refines the reference trajectory.

2.4 THEORETICAL ANALYSIS

We further analyze why historical prediction errors can serve as transferable knowledge for trajectory correction, focusing on phase-aligned trajectory errors and cross-cycle residual dynamics.

Transferability of Phase-Aligned Errors. Let $\mathcal{E}_h(\cdot)$ denote the residual function of the fixed base predictor at the h-th prediction step, such that

$$\mathbf{r}_{t+h} = \mathcal{E}_h(\mathbf{x}^{\text{ref}}_{t+h}). \tag{12}$$

We assume that $\mathcal{E}_h(\cdot)$ is locally smooth with respect to the phase representation:

$$\|\mathcal{E}_h(\mathbf{x}) - \mathcal{E}_h(\mathbf{x}')\|_2 \le L_{\mathcal{E}} \|\boldsymbol{\phi}(\mathbf{x}) - \boldsymbol{\phi}(\mathbf{x}')\|_2, \tag{13}$$

where $L_{\mathcal{E}} > 0$ is a local smoothness constant. Accordingly, for a retrieved historical case $i \in \mathcal{N}_{t,h}$,

$$\|\mathbf{r}_{t+h} - \mathbf{r}_{i+h}\|_2 \le L_{\mathcal{E}} \|\boldsymbol{\phi}_{t+h} - \boldsymbol{\phi}_{i+h}\|_2. \tag{14}$$

Under this smoothness assumption, phase-aligned matching thus favors historical cases with smaller upper bounds on residual discrepancy, supporting the transferability of their trajectory errors.

Transferability of Residual Dynamics. Phase-aligned errors may nevertheless evolve across orbital cycles. To characterize this evolution, we model the residual as the sum of a recurring phase-dependent component and a cycle-dependent component:

$$\mathbf{r}_{t+h} = \mathbf{b}_h(\boldsymbol{\phi}_{t+h}) + \mathbf{d}_{t+h}, \tag{15}$$

where $\mathbf{b}_h(\cdot)$ captures the recurring phase-dependent bias and $\mathbf{d}_{t+h}$ captures its cycle-dependent variation. For the phase-aligned predecessor $\pi(t,h)$,

$$\Delta\mathbf{r}_{t+h} = \mathbf{b}_h(\boldsymbol{\phi}_{t+h}) - \mathbf{b}_h(\boldsymbol{\phi}_{\pi(t,h)+h}) + \mathbf{d}_{t+h} - \mathbf{d}_{\pi(t,h)+h}. \tag{16}$$

Assuming that $\mathbf{b}_h(\cdot)$ varies smoothly with orbital phase, the phase-aligned predecessor minimizes the phase discrepancy within the candidate cycle range. When the resulting phase discrepancy is small, we have

$$\Delta\mathbf{r}_{t+h} \approx \mathbf{d}_{t+h} - \mathbf{d}_{\pi(t,h)+h}. \tag{17}$$

As a result, cross-cycle differencing suppresses recurring phase-dependent bias and isolates the cycle-to-cycle evolution of prediction errors. Finally, recall that $\hat{\mathbf{r}}_{t+h} = \mathbf{r}_{\pi(t,h)+h} + \sum_{i\in\mathcal{N}_{t,h}} \alpha_{t,h,i}\Delta\mathbf{r}_{i+h}$. Since $\mathbf{r}_{t+h} = \mathbf{r}_{\pi(t,h)+h} + \Delta\mathbf{r}_{t+h}$ and the routing weights form a convex combination,

$$\|\mathbf{r}_{t+h} - \hat{\mathbf{r}}_{t+h}\|_2 \le \sum_{i\in\mathcal{N}_{t,h}} \alpha_{t,h,i} \|\Delta\mathbf{r}_{t+h} - \Delta\mathbf{r}_{i+h}\|_2. \tag{18}$$

This bound shows that better alignment between transferred historical dynamics and the current cross-cycle evolution yields a tighter upper bound on the correction error, motivating adaptive residual-dynamics routing. The detailed proof is provided in Appendix B.

3 EXPERIMENTS

In this section, we empirically investigate four research questions:

- **RQ1 (Superiority and Efficiency):** Compared with representative baseline methods, can ORBITECHO achieve superior forecasting accuracy while remaining computationally efficient? (*c.f.* Tables 1 and 2)

Table 1: Orbit prediction results on three satellites (§ 3.2).

Method	Sentinel-3A		Sentinel-3B		SARAL		*Avg.*	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
SGP4*	3.92	5.57	3.92	5.57	3.96	5.63	3.94	5.59
Numerical Propagator*	0.13	0.15	0.17	0.25	0.24	0.28	0.18	0.23
TimesNet [ICLR23]	25.19	28.14	26.07	29.20	35.42	43.62	28.89	33.65
DLinear [AAAI23]	8.17	8.91	8.24	9.01	10.08	11.14	8.83	9.69
PatchTST [ICLR23]	619.47	620.84	619.45	620.82	589.33	590.54	609.42	610.73
iTransformer [ICLR24]	35.08	37.88	35.20	37.96	59.43	62.54	43.24	46.13
TimesFM [ICML24]	165.70	192.63	164.78	191.29	163.99	191.16	164.82	191.69
TimeMixer++ [ICLR25]	177.94	184.15	177.90	184.17	65.56	73.15	140.47	147.15
TimeMoE [ICLR25]	74.54	76.67	74.60	76.09	71.93	74.19	73.69	75.65
XLinear [AAAI26]	30.30	32.26	30.42	32.45	46.86	51.60	35.86	38.77
GTR [ICLR26]	48.10	51.46	48.12	51.55	73.78	76.91	56.67	59.97
DASR [KDD24]	68.99	76.16	69.09	76.17	90.42	100.49	76.16	84.27
AutoOrbit [KDD26]	4.21	19.56	4.20	19.55	16.80	27.56	8.40	22.22
ORBITECHO (Ours)	0.29	0.35	0.29	0.35	0.34	0.41	0.31	0.37

* SGP4 is propagated from the corresponding TLE records; Numerical Propagator performs high-fidelity propagation with perturbation models implemented in Orekit (CS-SI Orekit Contributors, 2026).

- **RQ2 (Generalization):** How well does ORBITECHO generalize to unseen satellites? (*c.f.* Table 3)
- **RQ3 (Component Effectiveness):** How do historical trajectory errors and residual dynamics contribute to its forecasting performance? (*c.f.* Table 4)
- **RQ4 (Robustness):** How well does ORBITECHO adapt to different forecasting backbones, extended prediction horizons, and challenging orbital conditions? (*c.f.* Tables 5 and 6, Figure 3)

3.1 EXPERIMENTAL SETUP

Data Curation. We evaluate ORBITECHO on five in-orbit satellites: Sentinel-3A, Sentinel-3B, SARAL, CryoSat-2, and Surface Water and Ocean Topography (SWOT), using precise orbit products from the International DORIS Service (IDS) (International DORIS Service, 2026) spanning 2021–2025. Each record contains position and velocity $[x, y, z, v_x, v_y, v_z]$ in the Geocentric Celestial Reference System (GCRS), sampled at one-minute intervals. To evaluate forecasting performance and cross-satellite generalization, we conduct experiments under two settings. i) *In-domain testing*, where Sentinel-3A, Sentinel-3B, and SARAL are split into $70\%/10\%/20\%$ for training, validation, and testing. ii) *Unseen-satellite evaluation*, where CryoSat-2 and SWOT are used only for evaluation. Dataset construction and preprocessing details are provided in Appendix C.1.

Baselines. We compare ORBITECHO with representative baselines from three categories. i) *Physics-based propagation*: SGP4 (Hoots & Roehrich, 1980) and a high-fidelity numerical propagator implemented with Orekit (CS-SI Orekit Contributors, 2026). ii) *General time-series forecasting models*: TimesNet (Wu et al., 2023), DLinear (Zeng et al., 2023), PatchTST (Nie et al., 2023), iTransformer (Liu et al., 2024), TimeMixer++ (Wang et al., 2025a), XLinear (Chen et al., 2026), GTR (Cao et al., 2026), and the large-scale time-series models TimesFM (Das et al., 2024) and Time-MoE (Shi et al., 2025). iii) *Orbit-specific models*: DASR (Jeong et al., 2024) and AutoOrbit (Zhang et al., 2026). Detailed descriptions of the baselines are provided in Appendix C.2.

Implementation Details. We use L=96 minutes of observations to predict the next H=60 minutes. Unless otherwise specified, DLinear (Zeng et al., 2023) serves as the base predictor and is trained using Adam (Kingma & Ba, 2015) with a learning rate of 1×10^{-4} and batch size 512. Afterward, the base predictor is fixed, and the residual correction module of ORBITECHO is optimized using AdamW (Loshchilov & Hutter, 2017) with a learning rate of 1×10^{-3} and batch size 128. Both stages are trained for up to 50 epochs. For historical error retrieval, we retain historical cases from the latest three orbital cycles and retrieve the top $K = 8$ phase-aligned cases. Forecasting accuracy is evaluated using the 3D position error in kilometers, defined as the Euclidean distance between predicted and ground-truth positions in GCRS, from which we report MAE and RMSE. Additional configurations are provided in Appendix C.4.

Table 3: Zero-shot evaluation on two unseen satellites (§ 3.3).

Method	CryoSat-2		SWOT		*Avg.*	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
SGP4	4.15	5.94	3.70	5.21	3.93	5.58
Numerical Propagator	0.14	0.16	0.11	0.15	0.12	0.15
TimesNet [ICLR23]	552.50	569.55	1245.88	1730.09	899.19	1149.82
DLinear [AAAI23]	23.22	27.39	34.53	36.47	28.88	31.93
PatchTST [ICLR23]	691.09	701.11	989.61	992.97	840.35	847.04
iTransformer [ICLR24]	162.48	171.71	168.33	183.01	165.41	177.36
TimesFM [ICML24]	153.92	182.18	169.52	195.19	161.72	188.69
TimeMixer++ [ICLR25]	447.34	456.93	794.61	808.55	620.98	632.74
TimeMoE [ICLR25]	72.73	77.29	76.75	78.49	74.74	77.89
XLinear [AAAI26]	241.23	273.42	294.08	329.59	267.65	301.50
GTR [ICLR26]	396.76	412.45	472.90	499.00	434.83	455.73
DASR [KDD24]	224.38	256.43	192.60	230.41	208.49	243.42
AutoOrbit [KDD26]	5845.37	9503.01	6968.72	11509.68	6407.05	10506.35
OrbitEcho (Ours)	0.33	0.39	0.28	0.33	0.30	0.36

3.2 Main Results (RQ1)

Performance Comparison. Table 1 presents the orbit forecasting performance of OrbitEcho against several representative baselines on three real-world satellites. As seen, OrbitEcho outperforms all learning-based baselines while remaining competitive with the physics-based numerical propagator. In particular, despite the complex architectures of modern forecasting models, *e.g.*, iTransformer, TimeMixer++, and XLinear, they exhibit limited prediction accuracy on orbit forecasting. A similar phenomenon is observed in large-scale time-series foundation models, *i.e.*, TimesFM and TimeMoE, which also yield suboptimal forecasting results. In addition, orbit-specific approaches (DASR and AutoOrbit) incorporate stronger inductive biases tailored to orbital dynamics, yet still underperform OrbitEcho.

Efficiency Analysis. Table 2 compares the computational efficiency of OrbitEcho with representative baselines on Sentinel-3A in terms of parameter count, GPU memory, and inference latency. As observed, OrbitEcho is lightweight, requiring only 0.16M parameters and 9.74 MB of GPU memory. Compared with orbit-specific methods, it uses up to 43.06× fewer parameters and 6.05× less GPU memory, while reducing inference latency by up to 1.44×. Together with its superior forecasting accuracy, these results demonstrate that OrbitEcho achieves a favorable trade-off between forecasting accuracy and computational efficiency.

Table 2: Computational efficiency on Sentinel-3A.

Method	Params.	GPU Mem.	Latency
DLinear	0.11M	8.16 MB	0.08 ms
PatchTST	9.77M	44.54 MB	0.65 ms
iTransformer	6.66M	33.58 MB	0.54 ms
TimesFM	203.57M	1035.72 MB	44.28 ms
XLinear	4.60M	11.57 MB	0.21 ms
DASR	0.18M	17.90 MB	0.85 ms
AutoOrbit	6.89M	58.97 MB	0.80 ms
OrbitEcho	0.16M	9.74 MB	0.59 ms

3.3 Cross-Satellite Generalization (RQ2)

We here evaluate the cross-satellite generalization of OrbitEcho on two unseen satellites, *i.e.*, CryoSat-2 and SWOT, without any fine-tuning. As reported in Table 3, general time-series forecasting models suffer from substantial performance degradation when applied to unseen satellites. In contrast, our OrbitEcho can generalize well to unseen satellites and consistently outperforms all learning-based baselines, attaining MAEs of only 0.33 km on CryoSat-2 and 0.28 km on SWOT.

3.4 Ablation Study (RQ3)

We evaluate the key designs of OrbitEcho on three satellites by progressively incorporating residual correction, phase-aligned residual matching, and residual dynamics modeling into the forecasting backbone ("Baseline"), as shown in Table 4.

Table 4: Ablation results of ORBITECHO on three satellites (§ 3.4).

Method	Sentinel-3A		Sentinel-3B		SARAL	
	MAE	RMSE	MAE	RMSE	MAE	RMSE
Baseline	8.17	8.91	8.24	9.01	10.08	11.14
+ Residual MLP	6.01	6.88	6.18	7.06	8.88	10.20
+ Historical Retrieval ($K = 1$)	0.48	0.56	0.49	0.57	0.58	0.68
+ Residual Averaging ($K = 8$)	0.34	0.40	0.33	0.41	0.62	0.76
+ Residual Dynamics (Full)	0.29	0.35	0.29	0.35	0.34	0.41

Table 5: Performance of ORBITECHO across different forecasting backbones (§ 3.5).

Satellite	Metric	PatchTST	XLinear	iTransformer	TimesFM
Sentinel-3A	MAE	0.53	0.43	0.73	0.13
	RMSE	0.60	0.51	0.82	0.17
Sentinel-3B	MAE	0.54	0.45	0.73	0.12
	RMSE	0.60	0.53	0.82	0.17
SARAL	MAE	6.58	0.48	0.78	1.16
	RMSE	7.73	0.57	0.87	1.60

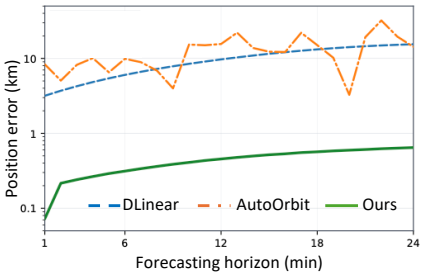


Figure 3: Prediction errors during a geomagnetic storm (§ 3.6).

Effect of Residual Correction. We first examine whether explicitly correcting prediction residuals can improve the reference forecast. To this end, we introduce a parameter-matched MLP corrector that directly predicts residual corrections from the reference trajectory, without exploiting historical prediction errors. As shown in Table 4, the Residual MLP consistently improves the baseline across all three satellites, reducing the MAE from 8.17 to 6.01 km on Sentinel-3A, from 8.24 to 6.18 km on Sentinel-3B, and from 10.08 to 8.88 km on SARAL.

Effect of Phase-Aligned Residual Matching. We next investigate whether historical trajectory errors provide more informative correction references than residual learning alone. By retrieving the most phase-aligned historical case ($K = 1$) and using its trajectory error for correction, the MAE is substantially reduced. Furthermore, averaging the trajectory errors from the top-8 phase-aligned cases further reduces the MAE to 0.34 and 0.33 km on Sentinel-3A and Sentinel-3B, while achieving 0.62 km on SARAL.

Effect of Residual Dynamics Modeling. Finally, we examine whether modeling how prediction errors change across adjacent orbital cycles can further improve trajectory correction. Compared with directly averaging the top-8 phase-aligned trajectory errors, ORBITECHO achieves lower errors across all three satellites, with the MAE on SARAL decreasing from 0.62 to 0.34 km. Overall, trajectory errors and residual dynamics provide complementary forms of transferable knowledge: the former provides correction references from recurring orbital states, while the latter captures their cross-cycle evolution for adaptation to the current prediction.

3.5 GENERALIZATION ACROSS FORECASTING BACKBONES (RQ4)

Table 5 evaluates ORBITECHO across forecasting backbones with diverse architectures and model scales. As shown, ORBITECHO achieves accurate predictions across all four backbones, attaining sub-kilometer MAEs in most cases despite their substantial architectural differences. Notably, XLinear and iTransformer maintain low prediction errors across all three satellites, while TimesFM achieves MAEs of 0.12–0.13 km on Sentinel-3A and Sentinel-3B.

3.6 ROBUSTNESS UNDER GEOMAGNETIC STORMS (RQ4)

We further evaluate the robustness of ORBITECHO on SARAL under challenging geomagnetic storm conditions, with the forecast initialized at 2025-06-02 11:02 UTC. Figure 3 compares the 3D position errors of ORBITECHO with DLinear and AutoOrbit. As shown, the prediction errors of DLin-

Table 6: Long-horizon orbit forecasting results on Sentinel-3A and Sentinel-3B (§ 3.8).

Method	Sentinel-3A						Sentinel-3B					
	2h		4h		6h		2h		4h		6h	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
DLinear	12.75	14.79	23.54	29.21	38.66	49.42	12.88	14.91	23.71	29.37	38.82	49.57
PatchTST	26.03	28.21	54.13	56.51	74.04	76.50	26.03	28.21	54.14	56.53	74.05	76.53
iTransformer	14.88	16.42	15.24	16.91	20.44	23.64	21.05	24.63	19.67	22.38	20.01	22.84
DASR	56.90	62.53	30.68	35.30	29.71	33.70	57.03	62.68	30.81	35.42	29.89	33.91
ORBITECHO	0.97	1.17	2.22	2.77	2.78	3.36	0.96	1.16	2.19	2.74	2.77	3.35

ear gradually increase with the forecasting horizon, while AutoOrbit exhibits large fluctuations and consistently high errors. In contrast, ORBITECHO maintains substantially lower and more stable prediction errors throughout the forecasting period, remaining below 1 km across all horizons.

3.7 CASE STUDY

To provide an intuitive view of how ORBITECHO corrects the reference trajectory, we present a qualitative case study on Sentinel-3A. As shown in Figure 4, the global view illustrates the orbital context and selected forecasting segment, while the local view compares the ground truth, reference prediction, and corrected prediction. The reference prediction progressively deviates from the ground-truth trajectory over the forecasting horizon, whereas ORBITECHO effectively corrects this deviation and remains closely aligned with the ground truth. Additional qualitative examples are provided in Appendix D.1.

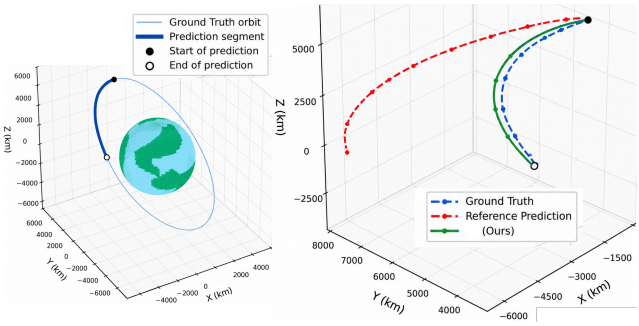


Figure 4: Qualitative visualization of trajectory correction on Sentinel-3A over a 24-minute forecasting horizon. Trajectory deviations are visually amplified for clarity.

3.8 LONG-HORIZON ORBIT FORECASTING (RQ4)

We further evaluate ORBITECHO over extended forecasting horizons of 2, 4, and 6 hours on Sentinel-3A and Sentinel-3B. As shown in Table 6, while the errors of existing methods generally increase with the forecasting horizon, ORBITECHO consistently achieves substantially lower errors on both satellites. Even at the 6-hour horizon, our ORBITECHO still maintains MAEs below 2.8 km, compared with over 20 km for all learning-based baselines, demonstrating its robustness to long-horizon orbit forecasting.

4 CONCLUSION

Conclusion. In this work, we introduced ORBITECHO, a residual correction framework that rethinks historical prediction errors as transferable knowledge for satellite orbit prediction. By transferring phase-aligned trajectory errors and their cross-cycle residual dynamics, ORBITECHO effectively leverages past prediction experience to refine current trajectories. Extensive experiments on real-world satellites demonstrate consistent improvements in forecasting accuracy, strong generalization across unseen satellites and robustness under severe geomagnetic storms.

Discussion. ORBITECHO relies on informative historical errors, which may be limited for newly observed satellites or orbital regimes with insufficient history. Moreover, it primarily exploits prediction experience without explicitly incorporating orbital physics. Future work could explore transfer under limited historical observations and integrate prediction experience with orbital physics for broader space-domain forecasting tasks.

Reproducibility Statement

To ensure the reproducibility of our results, we provide detailed descriptions of dataset construction and preprocessing, baseline configurations, and implementation details in the appendix. Complete proofs of the theoretical results are also provided. All datasets used in our experiments are obtained from publicly available sources, with corresponding links provided in the paper. We will release the implementation of OrbitEcho and experiment scripts upon publication.

AI Use Statement

Large language models (LLMs) were used solely for language editing, including improving the grammar, clarity, and style of the manuscript and suggesting alternative phrasings. All scientific ideas, methodological designs, experiments, analyses, and conclusions were developed independently by the authors. The authors reviewed and verified all LLM-assisted text and take full responsibility for the manuscript.

References

Francisco Caldas and Cláudia Soares. Machine learning in orbit estimation: A survey. *Acta Astronautica*, 220:97–107, 2024a.

Francisco Caldas and Cláudia Soares. Precise and efficient orbit prediction in LEO with machine learning using exogenous variables. In *2024 IEEE Congress on Evolutionary Computation (CEC)*, pp. 1–8, Piscataway, NJ, USA, 2024b. IEEE.

Fanpu Cao, Lu Dai, Jindong Han, and Hui Xiong. Enhancing multivariate time series forecasting with global temporal retrieval. *Proceedings of the International Conference on Learning Representations*, 2026.

Xinyang Chen, Huidong Jin, Yu Huang, and Zaiwen Feng. XLinear: A lightweight and accurate MLP-based model for long-term time series forecasting with exogenous inputs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2026.

Zhuqing Chen, Jiasen Yang, Zhongkang Yin, Yijia Wu, Lei Zhong, Qingyu Jia, and Zhimin Chen. Hybrid frequency-temporal modeling with transformer for long-term satellite telemetry prediction. *Applied Sciences*, 15(21):11585, 2025.

Qinghao Chu, Zhelong Wang, Yu Jiang, Pengrong Hou, Ruicheng Nie, Xin Shi, Fang Lin, Yuntong Kang, and Luchang Guo. KiGRU for long-term orbital prediction with kolmogorov-arnold networks. In *2025 IEEE International Conference on Systems, Man, and Cybernetics*, pp. 175–180, Piscataway, NJ, USA, 2025. IEEE.

CS-SI Orekit Contributors. CS-SI/Orekit: 13.1.4, 2026. URL https://doi.org/10.5281/zenodo.18527188. Version 13.1.4.

Abhimanyu Das, Weihao Kong, Rajat Sen, and Yichen Zhou. A decoder-only foundation model for time-series forecasting. In *Proceedings of the International Conference on Machine Learning*, 2024.

European Space Agency. ESA Space Environment Report 2025. https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025, 2025. Accessed: 2026-07-12.

Guangteng Fan, Xi Chen, Zhaoyue Chen, Ruichen Zhang, Peng Wu, Qihui Wei, Wenjun Xu, Jincheng Dai, and Lu Cao. Toward massive satellite signals of opportunity positioning: Challenges, methods, and experiments. *Space: Science & Technology*, 4:0191, 2024.

Wenhao Fan, Qingcheng Meng, Guan Wang, Hengwei Bian, Yabin Liu, and Yuanan Liu. Satellite edge intelligence: DRL-based resource management for task inference in LEO-based satellite-ground collaborative networks. *IEEE Transactions on Mobile Computing*, 24:10710–10728, 2025.

Felix R. Hoots and Ronald L. Roehrich. Models for propagation of NORAD element sets. Technical Report Project Space Track Report No. 3, Aerospace Defense Center, 1980.

International DORIS Service. DORIS Data and Products. https://ids-doris.org/data-products/tables-of-data-products.html, 2026. Accessed August 25, 2026.

SeungWon Jeong, Soyeon Woo, Daewon Chung, Simon S Woo, and Youjin Shin. Decomposed attention segment recurrent neural network for orbit prediction. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 5172–5182, New York, NY, USA, 2024. Association for Computing Machinery.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*, 2015.

Jae Yoo Lee and Seong Joon Yoo. Spatio temporal approach of transformer-based satellite orbit determination model. In *2025 International Conference on Artificial Intelligence in Information and Communication*, pp. 729–733, Piscataway, NJ, USA, 2025. IEEE.

Jiahao Liao, Hao Xu, Yufei Luo, and Yunhe Meng. A patch-transformer-based method for orbit prediction of LEO satellites. *Aerospace Science and Technology*, pp. 112652, 2026.

Richard J Licata and Piyush M Mehta. Uncertainty quantification techniques for data-driven space weather modeling: Thermospheric density application. *Scientific Reports*, 12(1):7256, 2022.

Chusen Lin and Junyu Chen. Using long short-term memory neural network for satellite orbit prediction based on two-line element data. *IEEE Transactions on Aerospace and Electronic Systems*, 61(4):8467–8475, 2025.

Chusen Lin, Junyu Chen, and Haoze Zhu. A method for compensating prediction errors to improve the orbit prediction accuracy of Starlink satellites. *GPS Solutions*, 29(1):17, 2025a.

Chusen Lin, Junyu Chen, Wei Liang, Bin Li, and Jizhang Sang. BiLSTM-driven error compensation for high-precision orbit prediction with two-line elements. *Space: Science & Technology*, 2026.

Shengsheng Lin, Weiwei Lin, Wentai Wu, Feiyu Zhao, Ruichao Mo, and Haotong Zhang. SegRNN: Segment recurrent neural network for long-term time series forecasting. *IEEE Internet of Things Journal*, 2025b. Early access.

Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. iTransformer: Inverted transformers are effective for time series forecasting. In *Proceedings of the International Conference on Learning Representations*, 2024.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.

Qingshan Luo, Ying Cao, Jiahao Ji, Yadi Song, Yurui Xu, Tao Yang, and Liangliang Xu. Improving satellite orbit prediction accuracy using a hybrid CNN-LSTM-transformer model. *CEAS Space Journal*, pp. 1–14, 2026.

Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *Proceedings of the International Conference on Learning Representations*, 2023.

Nathan Reiland, Aaron J Rosengren, Renu Malhotra, and Claudio Bombardelli. Assessing and minimizing collisions in satellite mega-constellations. *Advances in Space Research*, 67(11):3755–3774, 2021.

Xiaoming Shi, Shiyu Wang, Yuqi Nie, Dianqi Li, Zhou Ye, Qingsong Wen, and Ming Jin. Time-MoE: Billion-scale time series foundation models with mixture of experts. In *Proceedings of the International Conference on Learning Representations*, 2025.

Youjin Shin, Eun-Ju Park, Simon S Woo, Okchul Jung, and Daewon Chung. Selective tensorized multi-layer LSTM for orbit prediction. In *Proceedings of the ACM International Conference on Information and Knowledge Management*, pp. 3495–3504, New York, NY, USA, 2022. Association for Computing Machinery.

Shiyu Wang, Jiawei Li, Xiaoming Shi, Zhou Ye, Baichuan Mo, Wenze Lin, Shengtong Ju, Zhixuan Chu, and Ming Jin. TimeMixer++: A general time series pattern machine for universal predictive analysis. In *Proceedings of the International Conference on Learning Representations*, volume 2025, pp. 16980–17016, 2025a.

Yonglong Wang, Hongbin Ma, Zijun Que, Kunkun Liu, and Lujie Xiong. FDLSTM: A satellite orbit prediction model with frequency-domain feature fusion. *Array*, 2025b. Article 100564.

Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. TimesNet: Temporal 2d-variation modeling for general time series analysis. In *Proceedings of the International Conference on Learning Representations*, 2023.

Jiaping Xiao, Lu Bai, and Mir Feroskhan. Hybrid neural orbit propagation under sparse space-based observations. *IEEE Transactions on Aerospace and Electronic Systems*, 2026.

Yue Yang, Zhiqiang Wu, Saiyu Qi, and Fan Ma. Velocity-coupled representation refinement for satellite orbit prediction. *arXiv preprint arXiv:2608.23728*, 2026.

Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series forecasting? In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 11121–11128, Palo Alto, CA, USA, 2023. AAAI Press.

Donghao Zhang, Siyuan Liang, Boyuan Yang, Tao Qi, Shangguang Wang, and Qing Li. AutoOrbit: Physics-informed satellite orbit prediction. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 12820–12831, 2026.

Yan Zhang, Yong Wang, Haotong Cao, Yihua Hu, Zhi Lin, Kang An, and Dong Li. Self-similar traffic prediction for LEO satellite networks based on LSTM. *IET Communications*, 2024. Early access.