aerospace

MDPI

Article

A Dynamics-Informed Non-Causal Deep Learning Framework for High-Precision SOP Positioning Using Low-Quality Data

Zhisen Wang [1], Hu Lu [1,2,*] and Zhiang Bian [1]

[1] School of Information and Navigation, Air Force Engineering University, Xi'an 710077, China; 2019031031@cauc.edu.cn (Z.W.); bianzzaa@foxmail.com (Z.B.)
[2] National Key Laboratory of Unmanned Aerial Vehicle Technology, Northwestern Polytechnical University, Xi'an 710077, China
* Correspondence: sdkmsdn@sina.com

Abstract

Low Earth Orbit (LEO) satellite signals of opportunity (SOP) provide a viable positioning alternative in GNSS (Global Navigation Satellite System)-denied environments, yet their accuracy is fundamentally constrained by the low-quality orbital data typically available, such as SGP4 (Simplified General Perturbations model 4) predictions derived from Two-Line Elements (TLEs). To address this limitation, this paper proposes a dynamics-informed non-causal deep learning framework that enhances low-quality orbital data into high-fidelity trajectories for accurate SOP positioning. The proposed Non-Causal Dynamics-Informed Representation Temporal Convolutional Network (Non-Causal DIR-TCN) integrates phase space reconstruction and a Temporal Convolutional Network to explicitly model the chaotic dynamics inherent in LEO orbits, while relaxing the causality constraints of standard temporal convolutions to utilize both past and future context from the available SGP4 stream. Experimental results demonstrate that the framework significantly reduces orbit estimation errors and accelerates model convergence. When applied to LEO-SOP positioning, it achieves approximately 20% improvement in 2D positioning accuracy compared to conventional SGP4-based methods. This work effectively bridges the gap between accessible low-precision orbital data and high-accuracy state estimation, advancing the practical deployment of opportunistic signals for resilient positioning in challenging environments.

Keywords: signals of opportunity (SOP); orbit estimation; positioning; chaos; phase space reconstruction; Temporal Convolutional Network

Check for updates

Academic Editor: Marco Sabatini

Received: 19 January 2026
Revised: 3 March 2026
Accepted: 11 March 2026
Published: 13 March 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Low Earth Orbit (LEO)-based signals of opportunity (SOP) have become essential components of resilient positioning, navigation, and timing (PNT) architectures. By repurposing signals from non-cooperative LEO satellites—originally intended for communications or Earth observation—this approach offers a robust alternative in GNSS-denied environments. Compared to terrestrial counterparts, LEO-SOPs provide three distinct advantages: global coverage, higher received signal strength due to lower altitudes, and rapid geometric updates resulting from high satellite velocities. These characteristics enable reliable positioning in challenging scenarios such as urban canyons and contested electronic warfare environments, serving as a critical fail-safe where conventional systems degrade [1–4].

However, employing LEO constellations for PNT confronts a fundamental barrier: these systems were not designed for navigation, primarily lacking precise real-time

ephemerides. Consequently, practitioners often rely on publicly available Two-Line Element (TLE) sets propagated via the SGP4 model. While analytical models like SGP4 are computationally efficient for real-time applications, they are fundamentally limited by simplified dynamical assumptions. Neglecting complex perturbations—such as atmospheric drag, solar radiation pressure, and non-spherical gravitational effects—results in kilometer-level position errors that accumulate over time [5–7]. Although high-precision numerical propagators (e.g., High-Precision Orbit Propagation, HPOP) exist, their computational intensity renders them unsuitable for real-time navigation in embedded receivers. Thus, the absence of accurate, real-time satellite states remains the predominant bottleneck in SOP-based positioning [8,9].

In response to these challenges, deep learning has emerged as a modern trend in addressing complex space-related tasks. As noted in recent reviews by Khoroshylov and Redka [10] and Silvestrini and Lavagna [11], data-driven approaches are increasingly leveraged for spacecraft guidance, navigation, and control (GNC), offering the capability to model complex dynamics that are difficult to capture with traditional analytical methods. Following this trend, deep learning techniques have been applied to refine empirical atmospheric density models for precise orbit prediction [12–14] and to determine the attitude of uncooperative debris via image analysis [12,15].

Physics-informed machine learning for orbital state estimation and prediction can be broadly categorized into four research directions: (i) PINNs and ODE-constrained networks that embed equations of motion as soft constraints in the loss to promote dynamical consistency [16]; (ii) hybrid residual learning approaches that treat a physics propagator (e.g., SGP4) as a baseline and learn a correction term between low-fidelity and high-fidelity or measured ephemerides [17]; (iii) differentiable or learned propagators that replace part of numerical/analytical propagation with trainable components while preserving physical structure [18]; and (iv) physics-guided representation learning that injects dynamical priors through carefully designed state representations, embeddings, or invariant/attractor-aware features, improving learning efficiency when high-fidelity labels are limited [19].

For LEO satellites and megaconstellations, recent studies have investigated learning-based corrections to TLE/SGP4 propagation, uncertainty estimation, and ephemeris error mitigation for navigation and space situational awareness applications. These results suggest that enhancing widely accessible but low-quality orbit products can substantially improve downstream positioning performance. In addition, PINN-based orbit determination has been explored by integrating physical laws directly into network architectures to improve robustness [20], and Transformer-based models have been proposed to capture spatio-temporal dependencies in orbital dynamics, enabling rapid predictions suitable for operational decision-making [21,22]. Collectively, these advances indicate a promising role for deep learning in augmenting traditional physics-based pipelines for space situational awareness [23].

Despite these achievements, existing learning-based orbit estimation studies still exhibit several gaps depending on the methodological family. Residual learning and purely data-driven predictors often learn historical prediction errors as targets without explicitly encoding dynamical system structure. In contrast, PINN-based approaches embed equations of motion as soft constraints but may require careful tuning and may not directly address the practical setting where only low-quality public orbit products (e.g., SGP4/TLE at high rate) are available as inputs for downstream navigation.

To address these challenges, we propose a dynamics-informed, non-causal deep learning framework. Our objectives are threefold: (1) to analyze satellite orbital dynamics and incorporate physical insights into state estimation; (2) to develop a framework capable of transforming imprecise SGP4 outputs into high-precision trajectories comparable to HPOP;

and (3) to enhance LEO-SOP positioning accuracy using these refined states. Specifically, our approach follows the hybrid residual learning/physics-guided representation direction: SGP4/TLE provides a readily available low-fidelity state stream, while the model learns a mapping toward an HPOP-equivalent trajectory. Rather than enforcing an explicit force model constraint as in PINNs, we incorporate dynamical system priors by applying phase space reconstruction as a physically motivated state encoder to expose underlying attractor geometry before temporal modeling, and by using Lyapunov exponents to quantify orbit chaos strength under different time-window lengths. We find that the estimated chaoticity weakens as the time window increases, indicating improved effective predictability with longer context; since TLEs are typically released every 24 h, we therefore adopt a 24 h (maximum) history window as model input to operate in the most stable correction regime.

Broadly, this framework also provides a generalizable methodology for enhancing SGP4 data, with potential applicability to other domains involving cost-effective sensing, such as environmental monitoring and remote sensing. The contributions and novelties are summarized as follows:

- We conducted a systematic dynamical analysis on a 6D LEO state time series $[x, y, z, v_x, v_y, v_z]$ generated by propagating public TLEs with SGP4 at 1 Hz over 24 h, including AMI/FNN-based embedding assessment and Lyapunov exponent estimation, and used these insights to motivate a non-causal estimation architecture.
- We propose a dynamics-informed encoder–decoder framework for satellite state estimation. In the encoder, we do not rely solely on neural networks. Instead, phase space reconstruction functions as the primary physical encoder, while neural networks provide auxiliary encoding. For the decoder, we enhanced the Temporal Convolutional Network (TCN) to better capture temporal contextual dependencies.
- We extensively validated the proposed framework in a real-world LEO-SOP positioning scenario. Experimental results show that our method substantially improves positioning accuracy relative to conventional approaches using raw SGP4 outputs. These results demonstrate the method's efficacy and applicability in realistic GNSS-denied environments.

2. Dynamics-Informed Encoding: Orbit Chaos Motivated Phase Space Reconstruction

In this work, "dynamics-informed" does not refer to enforcing the equations of motion as hard constraints (e.g., PINNs). Instead, we inject orbit dynamics knowledge into the representation by reconstructing a phase space embedding of the ephemeris time series. The key idea is that LEO ephemerides exhibit strong nonlinear dynamics and sensitivity to perturbations; therefore, the mapping from low-fidelity SGP4 trajectories to high-fidelity reference trajectories benefits from a representation that preserves the underlying dynamical state. Consequently, we perform time delay embedding on the six-dimensional ephemeris (x, y, z, v_x, v_y, v_z) and use the reconstructed phase space as a dynamics-informed encoder to provide richer dynamical context while avoiding redundant raw inputs.

Compared with a purely data-driven encoder that learns latent features directly from raw sequences, our dynamics-informed encoder is explicitly constructed to preserve the dynamical-state structure (attractor geometry) implied by the ephemeris time series. This design reduces the burden on the neural network to "rediscover" basic dynamical relationships from limited high-fidelity labels and improves learning efficiency under practical data constraints.

We further analyze the chaotic characteristics of LEO ephemerides by computing Lyapunov exponents under different time-span settings (Section 2.3). Considering the typical ∼24 h TLE update cycle, we adopt a 24 h window (86,400 s) as the sequence length

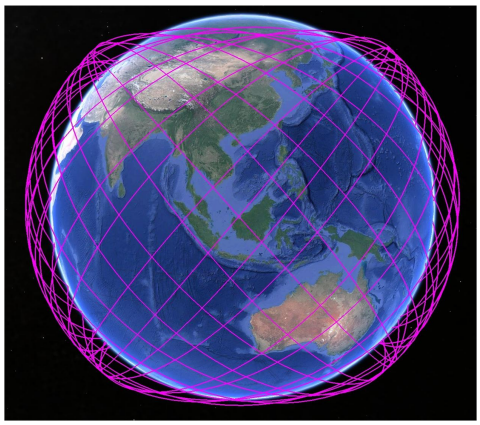
to align with the data availability while maximizing dynamical observability and mitigating chaos-induced uncertainty.

In light of the characteristics of ephemeris data $[x, y, z, v_x, v_y, v_z]$, our methods for chaos characteristics encompass the following aspects:

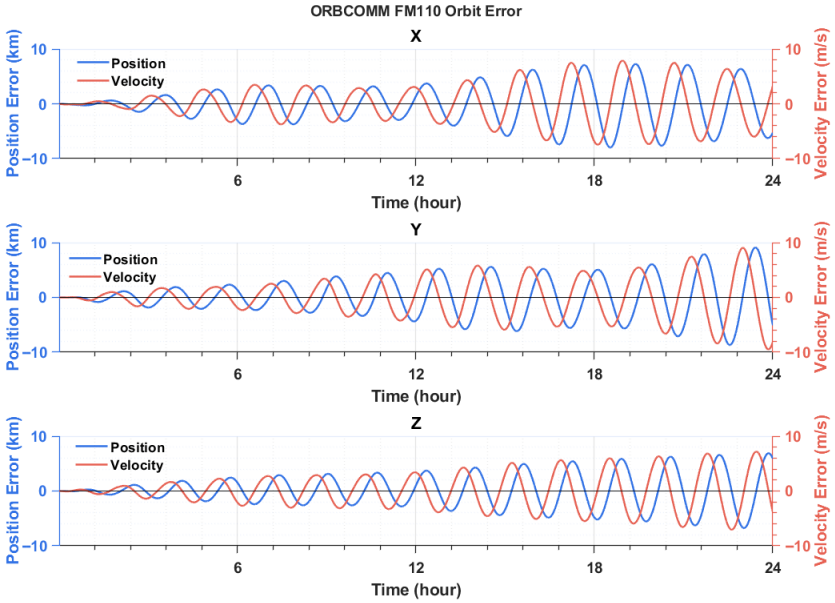
1. Phase space reconstruction is employed to recover the chaotic dynamics.
2. The calculation of Lyapunov exponents is carried out to quantify the chaotic behavior in the orbital motion.

2.1. Phase Space Reconstruction

As chaotic behavior pervades numerous complex systems, characterizing their chaotic features from acquired time series data constitutes a critical research priority. The satellite trajectory and time series data are shown in Figure 1.



(**a**) 24-hour satellite trajectory



(**b**) LEO satellite trajectory data

Figure 1. Visualization of the complexity in LEO satellite orbital dynamics. (**a**) The 3D trajectory of the target satellite (ORBCOMM FM110) over 24 h. (**b**) Orbit error (residuals) analysis of ORBCOMM FM110.

Successful phase space reconstruction via time-delayed embedding requires the determination of two critical parameters: the delay interval (τ), defining the temporal offset for

state-space unfolding, and the embedding dimension (m), specifying the reconstruction space dimensionality. These parameters enable approximation of original system dynamics from multivariate time series data [24].

For empirical time series, both parameters typically require estimation as they are seldom known a priori. Established techniques for scalar time series include computing two functions: 1. Average Mutual Information (AMI)—where its initial local minimum indicates optimal τ. 2. False Nearest Neighbors (FNN)—where convergence onset determines embedding dimension m.

This study employs an enhanced phase space reconstruction framework [24], extending both methods to multivariate time series for joint estimation of τ and m. For the delay parameter τ, we implement uniform multivariate average mutual information. Regarding embedding dimension m, we use a straightforward extension of the original false nearest neighbor's algorithm by Kennel et al. [25].

2.1.1. Average Mutual Information

Fraser and Swinney [26] established a method for determining optimal embedding delays in scalar time series by maximizing statistical independence between coordinates. Their approach quantifies dependence through mutual information $I(x_t, x_{t+\tau})$ between the original series x_t and its τ-lagged counterpart. This self-referential computation constitutes the average mutual information (AMI) metric. Serving as a nonlinear analog to autocorrelation, AMI is mathematically defined as:

$$I(x(t), x(t+\tau)) = \sum_{i,j} p_{ij}(\tau) \log\left(\frac{p_{ij}(\tau)}{p_i p_j}\right) \tag{1}$$

Here, p_i denotes the probability of $x(t)$ occupying the i-th histogram bin derived from empirical data, while $p_{ij}(\tau)$ represents the joint probability of $x(t)$ residing in bin i and $x(t+\tau)$ in bin j. Crucially, only the joint probability $p_{ij}(\tau)$ exhibits τ-dependence. Furthermore, AMI function values are governed by histogram construction parameters—specifically bin width and alignment. As illustrated in Figure 2, we summarize the procedure for calculating the average mutual information as follows:

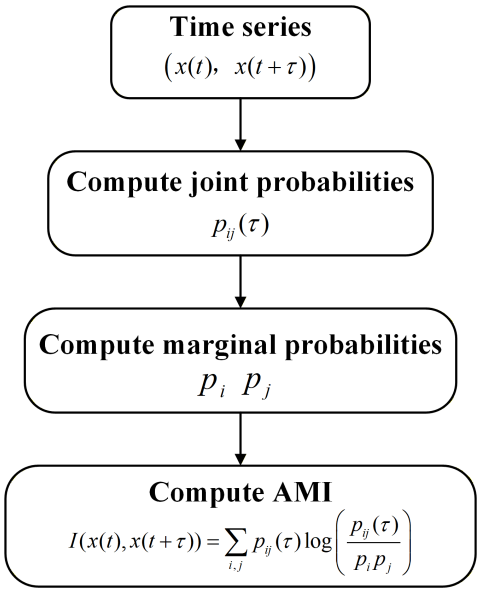


Figure 2. The calculation procedure of the average mutual information method.

Fraser and Swinney [26] postulated selecting the first minimum of $I(x(t), x(t+\tau))$ as the optimal τ for generating maximally decorrelated coordinates in time delay embedding. This specific τ-value ensures the primary coordinate $y_1(t) = x(t)$ maintains statistical independence from subsequent coordinates ($y_2(t) = x(t+\tau)$, etc). When AMI exhibits monotonic decay rather than local minima—a frequent empirical occurrence—alternative criteria apply, such as the smallest τ where AMI falls below $1/e$ of its initial magnitude.

To obtain maximally independent coordinates for time-delayed phase space embedding, we select the first minimum of $I(x(t), x(t+\tau))$ as the optimal τ value. This specific choice ensures that the first coordinate $y_1(t) = x(t)$ achieves maximal independence from the second coordinate $y_2(t) = x(t+\tau)$ and subsequent dimensions.

Within this method, we estimate a unified time delay τ for multivariate time series by computing the average mutual information (AMI) across all data dimensions, yielding a unified delay parameter.

2.1.2. False Nearest Neighbors

As shown in Figure 3, the false nearest neighbors algorithm is a widely used method for determining the minimal embedding dimension m required to reconstruct the phase space of a dynamical system from a time series. The core idea behind FNN is to identify whether neighboring points in a lower-dimensional embedding remain neighbors when projected into a higher-dimensional space. If the distance between neighbors changes significantly upon embedding, they are classified as "false neighbors," indicating that the current embedding dimension is insufficient to capture the system's true dynamics.

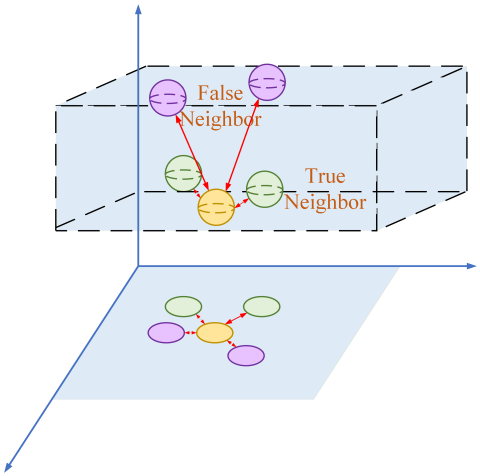


Figure 3. False nearest neighbors.

Firstly, we consider a scalar time series. Takens' theorem provides the foundation for reconstructing this latent phase space through time delay embedding. This involves constructing an m-dimensional trajectory vector $\mathbf{y}(t)$ from the original series $x(t)$:

$$\mathbf{y}(t) = \left[x(t),\, x(t+\tau),\, \ldots,\, x\big(t+(m-1)\tau\big)\right] \tag{2}$$

where t and τ are the time index and delay, respectively, and m represents the embedding dimension.

To determine the minimal m, we analyze the behavior of neighboring points. Let $\mathbf{y}^{(r)}(t)$ be the r-th nearest neighbor of $\mathbf{y}(t)$ in the m-dimensional space. Their Euclidean distance $R_m(t,r)$ is:

$$R_m^2(t,r) = \sum_{k=0}^{m-1} \left[x(t+k\tau) - x^{(r)}(t+k\tau)\right]^2 \tag{3}$$

where $x^{(r)}(t + k\tau)$ denotes the $(k + 1)$-th component of the neighbor vector $\mathbf{y}^{(r)}(t)$. When the embedding dimension increases to $m + 1$, the new distance is:

$$R^2_{m+1}(t, r) = R^2_m(t, r) + \left[x(t + m\tau) - x^{(r)}(t + m\tau)\right]^2 \quad (4)$$

The key insight is this: If the current embedding dimension m is sufficient to unfold the system's attractor, adding another dimension should not substantially alter the relative proximity of true neighbors. If, however, the distance increases dramatically (indicating the points were spuriously close in the lower-dimensional projection), they are labeled false neighbors. This signifies that m is inadequate.

The proportion of false nearest neighbors (FNN) is calculated for each embedding level. An adequate m is indicated when this proportion drops to near zero (or stabilizes at a minimal level) upon increasing m, signifying geometric consistency has been achieved.

2.1.3. Extension to Multidimensional Time Series

Let $x(t) \in \mathbb{R}^p$ be the p-dimensional state, m be the embedding dimension (number of delays), and $M = mp$ the total reconstructed dimension.

We extend this logic to p-dimensional input data $\mathbf{x}(t) = [x_1(t), x_2(t), \ldots, x_p(t)]$. Increasing the embedding count m by 1 corresponds to augmenting the phase space by p dimensions. The squared distance in the current space (with total dimension $M = m \times p$) is:

$$R^2_M(t, r) = \sum_{j=1}^{p} \sum_{k=0}^{m-1} \left[x_j(t + k\tau) - x_j^{(r)}(t + k\tau)\right]^2 \quad (5)$$

To identify "false neighbors," we employ two criteria established by Kennel et al. [25]:

1. Distance Ratio:

$$\sqrt{\frac{R^2_{M+p}(t, r) - R^2_M(t, r)}{R^2_M(t, r)}} > R_{\text{tol}}$$

where R_{tol} is typically set to 15.

2. Absolute Distance:

$$\frac{R_{M+p}}{\sigma_y} > A_{\text{tol}},$$

where σ_y is the standard deviation of the data and $A_{\text{tol}} \approx 2$.

If either criterion is met, the neighbor is "false." The optimal m is reached when the proportion of false nearest neighbors (FNN) drops to near zero, indicating the attractor is fully unfolded. The final reconstruction dimension is $M = m \times p$.

Increasing the embedding count to $m + 1$ adds the vector $x(t + m\tau)$ to the trajectory description, yielding a $M + p$-dimensional space. The new squared distance is:

$$R^2_{M+p}(t, r) = R^2_M(t, r) + \sum_{j=1}^{p} \left[x_j(t + m\tau) - x_j^{(r)}(t + m\tau)\right]^2 \quad (6)$$

The FNN proportion is computed for successive values of m. We select the smallest m for which the FNN proportion becomes negligible (or stabilizes). The corresponding reconstructed dimension is $M = mp$.

2.2. Computation of Lyapunov Exponents

This section employs an enhanced version of Rosenstein's [27] algorithm, as proposed by Mehdizadeh [28], for Lyapunov exponent computation.

To enhance the robustness of Lyapunov exponent (LyE) estimation in noisy signals, we use a modification to the Rosenstein et al. [27] algorithm. The key revision involves using multiple nearest neighbors instead of a single point to compute the initial separation distance, thereby reducing the influence of noise-induced false neighbors.

The initial average separation $d_{j,\mathrm{ave}}(0)$ is calculated as the mean Euclidean distance between the reference point X_j and its n nearest neighbors:

$$d_{j,\mathrm{ave}}(0) = \frac{1}{n}\sum_{k=1}^{n} \|X_{jk} - X_j\| \tag{7}$$

where $\|\|$ denotes the Euclidean norm, and X_{jk} represents the k-th nearest neighbor of X_j. The neighbors are selected such that their distances exceed the mean period of the time series to avoid temporal correlation.

Subsequently, the algorithm follows the standard Rosenstein algorithm. The largest finite-time Lyapunov exponent $\lambda_{\max}$ is derived from the logarithmic divergence of neighboring trajectories:

$$d(t) = d_{\mathrm{ave}}(0) \cdot e^{\lambda_{\max} t} \tag{8}$$

For each reference–neighbor pair, the logarithmic divergence at time $i\Delta t$ is approximated as:

$$\ln\left[d_{j,\mathrm{ave}}(i)\right] \approx \ln[d_{\mathrm{ave}}(0)] + \lambda_{\max}(i\Delta t) \tag{9}$$

To determine the maximum Lyapunov exponent $\lambda_{\max}$, let $y(i)$ denote the average logarithmic divergence at the i-th discrete time step:

$$y(i) = \left\langle \ln d_{j,\mathrm{ave}}(i) \right\rangle_j, \tag{10}$$

where $\langle\cdot\rangle_j$ denotes the average over all reference points (and their neighbor sets) indexed by j, and $\ln(\cdot)$ is the natural logarithm.

$$y(i) \approx \lambda_{\max} \cdot (i \cdot \Delta t) + C \tag{11}$$

where Δt is the sampling period and C is a constant.

To define the regression process, $\lambda_{\max}$ is estimated using the ordinary least-squares (OLS) method over the linear growth region $[i_1, i_2]$ as follows:

$$\lambda_{\max} = \frac{1}{\Delta t} \frac{\sum_{i=i_1}^{i_2} (i - \bar{i})(y(i) - \bar{y})}{\sum_{i=i_1}^{i_2} (i - \bar{i})^2} \tag{12}$$

where $\bar{i}$ and $\bar{y}$ are the arithmetic means of i and $y(i)$ within the specified range, respectively.

In our implementation, the sampling frequency is $f_s = 1$ Hz, the dominant period is set to $T = 600$ s, and the divergence curve is tracked over a window of length $w_s = 3600$ s (i.e., $w_s f_s = 3600$ samples). Following the least-squares fitting procedure in the Rosenstein-type algorithm, we estimate two finite-time slopes from the averaged logarithmic divergence $y(i)$: (i) the short-term exponent λ_s is obtained by fitting $y(i)$ over the first half period, $i \in [1,\ 0.5\,T f_s] = [1,\ 300]$, $t \in [1,\ 300]$ s, and (ii) the long-term exponent λ_l is obtained by fitting $y(i)$ over the interval from four periods up to the window end, $i \in [4\,T f_s,\ w_s f_s] = [2400,\ 3600]$, $t \in [2400,\ 3600]$ s.

This modification mitigates noise sensitivity by leveraging multiple neighbors, ensuring more reliable LyE estimation while preserving the original method's theoretical framework.

2.3. Analysis of Chaotic Characteristics

In the following analysis, we apply the aforementioned phase space reconstruction methodology to ephemeris data $[x, y, z, v_x, v_y, v_z]$, subsequently computing the corresponding Lyapunov exponents to characterize system dynamics.

Time delay parameter estimation was performed on all six ephemeris dimensions (x, y, z, v_x, v_y, v_z). The upper bound for time delays can be constrained via parameter adjustment, with optimal values determined through iterative evaluation of output results.

As illustrated in Figure 4a, the AMI functions exhibit local minima at τ = 575, 613, 39, 649, 536, and 41 for the respective dimensions. We compute the arithmetic mean of these dimension-specific τ estimates, yielding $\tau \approx 409$. Consequently, we adopt $\tau = 409$ as the unified integer-valued embedding delay. It should be noted that the data utilized for these parameter estimations is the normalized SGP4 orbital data.

The estimation of τ for the six-dimensional state reveals varying dynamics across different coordinates. Specifically, while the X and Y components exhibit local minima at larger lags, the Z-axis components reach their first minima earlier, likely due to the higher sensitivity of cross-track perturbations or data noise. To ensure a balanced representation for the multivariate encoder, we utilize the arithmetic mean of these values to define a unified embedding delay $\tau = 409$. This strategy, combined with the subsequent feature extraction by the encoder's fully connected layers, ensures that the model remains robust to dimension-specific noise while capturing the global attractor geometry.

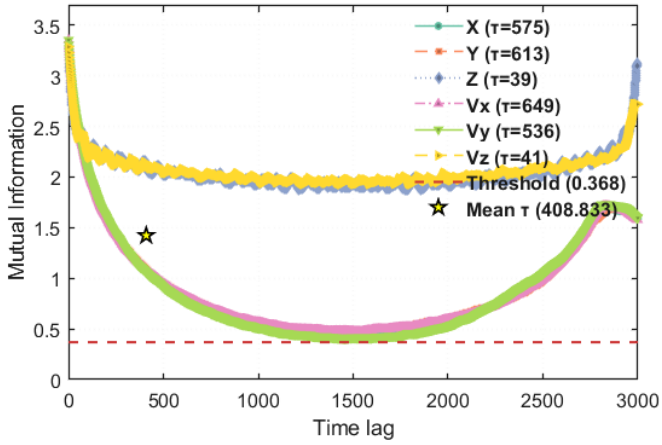
We analyze the six-dimensional ephemeris with data structured in column-wise dimensions (rows representing data points), using delay parameter $\tau = 409$. The FNN function plotted in Figure 4b demonstrates that the original 6D time series already possesses sufficient intrinsic dimensionality.

Following phase space reconstruction, we compute the Lyapunov exponents (LyE) using the previously described methodology to analyze the orbital dynamics. The green line in Figure 5a represents the divergence curve, quantifying the average separation evolution of adjacent trajectories in state space. The orange line denotes the short-term Lyapunov exponent (LyEs = 2.3595×10^{-7}), derived from the initial linear segment's slope and characterizing local dynamic stability. The blue line indicates the long-term Lyapunov exponent (LyEl= 2.5706×10^{-7}), which reflects the mean divergence rate over extended evolution and serves as the primary indicator of global chaotic characteristics.

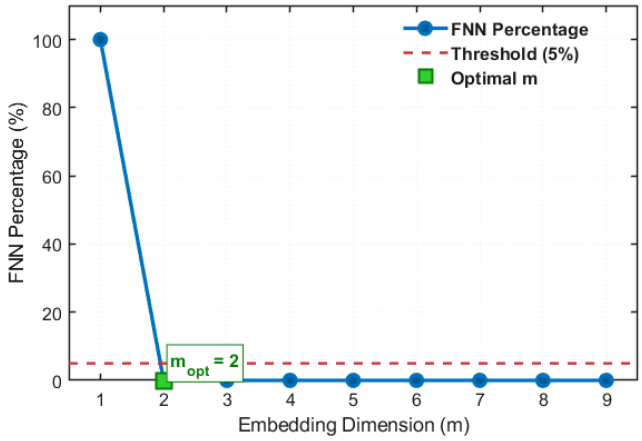
Both LyE_s and LyE_l exhibit positive values ($\mathrm{LyE}_s > \mathrm{LyE}_l > 0$), confirming chaotic behavior in the system. Although weak, this chaos manifests as nonlinearity and instability in satellites motion. Long-term dynamics are constrained by gravitational forces and orbital design, suppressing chaos, whereas short-term evolution exhibits heightened sensitivity to initial condition perturbations (e.g., launch errors) and environmental disturbances. This enhanced uncertainty complicates short-term prediction.

Moreover, our investigation reveals that the length of the input time series exerts a discernible influence on the strength of chaotic characteristics, as measured by the Lyapunov exponent. As illustrated in Figure 5b, shorter sequences tend to amplify chaotic behavior, whereas extending the observation window attenuates it. This phenomenon may be attributed to the gradual reduction in orbital uncertainty as the satellite follows its deterministic trajectory. Consequently, to minimize the impact of chaos, it is advisable to employ the longest feasible time series. Taking into account the approximately 24 h update cycle of TLE data, we selected a duration of 86,400 s (24 h) for the time series length. This duration aligns with the TLE publication interval while maximizing observational continuity.

Applications such as opportunistic signal positioning demand high-precision, real-time satellite state data. Overcoming prediction uncertainties induced by chaotic characteristics is therefore essential for achieving precise LEO satellite state forecasting.

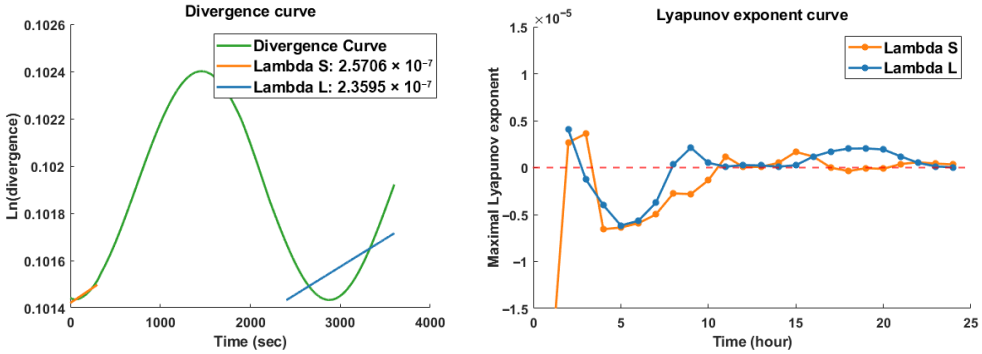


(**a**) Time delay parameter τ



(**b**) Embedded dimension parameter m

Figure 4. Determination of phase space parameters. (**a**) The calculation result of the time delay parameter τ. (**b**) The calculation result of the embedded dimension parameter.



(**a**) Chaotic Characteristics Analysis (**b**) Variation Trend of Chaotic Characteristics

Figure 5. Orbital dynamics analysis. (**a**) Computing the chaotic characteristics of a satellite orbit using 24 h data as input. (**b**) Calculating the variation in orbital chaotic characteristics using data of different time spans.

3. Orbit Uncertainty Estimation Framework

Our goal is to learn a supervised mapping that upgrades a readily available low-fidelity orbit state stream produced by SGP4/TLE to a high-fidelity reference trajectory produced by HPOP over the same time interval [29–36]. The input is a 6D state sequence

$\mathbf{s}^{\mathrm{SGP4}}(t) = [x, y, z, v_x, v_y, v_z]^\top$, and the target is the corresponding HPOP state sequence $\mathbf{s}^{\mathrm{HPOP}}(t)$ sampled at the same rate. Therefore, the problem is a regression (sequence-to-sequence alignment) task rather than a forecasting task [37,38].

Importantly, in our application, the full SGP4 sequence of a time segment is available once the TLE is obtained and propagated, so using a limited symmetric temporal context around each time step does not violate any physical causality. This motivates a non-causal temporal model: for estimating the state at time t, the network is allowed to use a bounded window $[t - \Delta,\ t + \Delta]$ from the SGP4 stream to counteract the short-term sensitivity indicated by the positive Lyapunov exponent (Section 2.2). The Lyapunov exponent analysis is used as physical motivation and diagnostic evidence for short-term sensitivity; it is not an explicit input feature nor a loss constraint in the current implementation.

The proposed framework utilizes an encoder–decoder architecture to map SGP4 inputs to high-fidelity HPOP-equivalent orbital states. The encoder integrates phase space reconstruction—acting as a physically-informed prior—with neural layers to extract salient dynamical features. For the decoder, we adapt the Temporal Convolutional Network (TCN) architecture by relaxing strict causality; this allows the model to leverage a symmetric temporal context (incorporating both past and limited future SGP4 information) to enhance state estimation.

3.1. Overall Architecture

This section presents an overview of the proposed orbital uncertainty estimation framework non-causal DIR-TCN (Non-Causal Dynamics-Informed Representation TCN, Non-Causal DIR-TCN). As illustrated in Figure 6, the framework uses an encoder–decoder architecture. The framework takes SGP4 data, generated from TLE files, as input. The next step is phase space reconstruction of the input time series $[x, y, z, v_x, v_y, v_z]$. A fully connected layer then encodes the reconstructed data to extract salient features. Finally, a modified Temporal Convolutional Network (TCN) serves as the decoder. A final fully connected layer maps the decoder output to the estimated satellite-state time series.

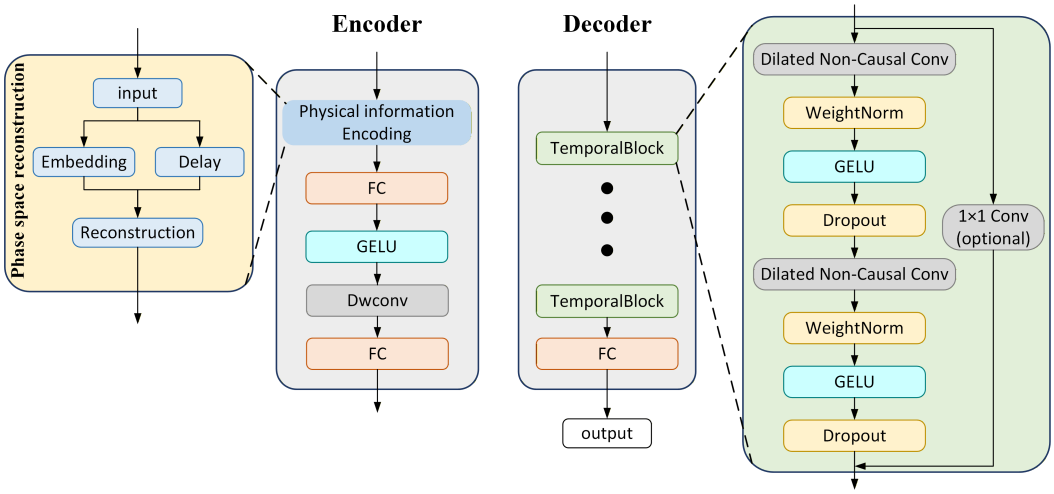


Figure 6. The architecture of non-causal DIR-TCN. Lyapunov analysis provides motivation for using non-causal context; it is not a computational block.

3.2. Encoder Block

Note that SGP4 provides a 6D state $[x, y, z, v_x, v_y, v_z]$, which can already be regarded as a phase space representation. In this paper, "phase space reconstruction" is used in a general dynamical systems sense as a time delay embedding framework to (i) estimate characteristic delay τ and embedding sufficiency via AMI/FNN, and (ii) provide a unified physically motivated encoder for time series modeling. When the FNN analysis indicates

that the original 6D series is already sufficient, the reconstruction step degenerates to a minimal embedding (i.e., no additional delay coordinates are appended), and the encoder passes the state sequence forward without increasing dimensionality.

To balance predictive power with computational efficiency, we implement this hybrid encoder by combining the aforementioned dynamics-informed embedding with a neural-based component. Instead of relying on a high-dimensional black-box architecture, the encoder leverages the reconstructed phase space to recover the underlying attractor's intrinsic structure. Within this framework, the dynamics-informed stage provides a principled representation—either as an augmented time delay sequence or as the original 6D state when sufficiency is met. Subsequently, a streamlined neural network layer extracts salient features, filtering redundant information and compressing the phase space data into an essential representation. This synergistic approach captures the fundamental dynamics while significantly reducing model complexity and computational overhead.

3.3. Decoder Block

3.3.1. Non-Causal Convolution

The original Temporal Convolutional Network (TCN) was built upon two foundational principles: the network must produce an output sequence of the same length as the input sequence, and no leakage of future information into past computations should occur. To adhere to the first principle, TCN employs a one-dimensional fully convolutional network (FCN) architecture. In this design, each hidden layer maintains the same length as the input layer. Zero padding of length (kernel size $-$ 1) is applied to ensure that subsequent layers preserve the temporal length of previous ones. To satisfy the second principle, TCN utilizes causal convolutions. This ensures that the output at time step t is convolved only with elements from time step t and earlier in the previous layer, thereby strictly preventing any access to future information.

However, considering the specific context of our task—where a segment of SGP4 data is already available and needs to be converted into corresponding HPOP data over the same period—strict adherence to temporal causality is unnecessary. Unlike conventional approaches that constrain predictions to rely exclusively on past observations, our method incorporates both past and future information to infer the current state. Specifically, we revise the padding computation method in causal convolution by replacing Equation (13) with Equation (14).

$$padding = (kernel_size - 1) * dilation \tag{13}$$

$$padding = (kernel_size - 1) * dilation/2 \tag{14}$$

where *dilation* represents dilated convolutions; we will introduce this in the next section.

3.3.2. Dilated Convolutions

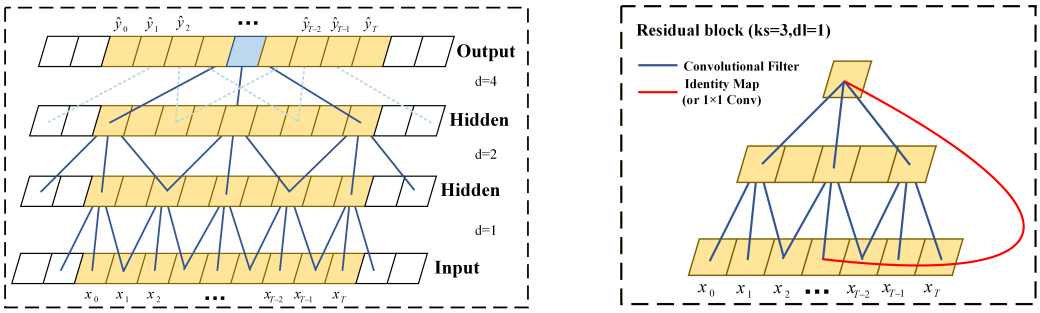
Solely utilizing non-causal convolutions for time series modeling presents a fundamental limitation: their receptive field grows only linearly with network depth. Consequently, they are ill-suited for tasks that require capturing long-range dependencies. To address this, we follow the methodology of TCN by adopting dilated convolutions, which achieve a significantly larger receptive field. For a 1-D sequence input $\mathbf{x} \in \mathbb{R}^n l$ and a filter $f : \{0, ..., k-1\} \rightarrow \mathbb{R}$, the dilated convolution F at position s is mathematically expressed as:

$$F(s) = (\mathbf{x} *_{dl} f)(s) = \sum_{i=0}^{k-1} f(i) \cdot \mathbf{x}_{s-dl*i} \tag{15}$$

In this formulation, dl denotes the dilation factor and ks is the filter size. This configuration effectively introduces a fixed interval between consecutive filter taps. Notably,

setting $dl = 1$ recovers a standard convolution. Progressively increasing the dilation factor allows a neuron at deeper layers to integrate information from a broader input range, thereby enlarging the network's receptive field exponentially.

Similar to the standard TCN, there exist two principal strategies for enlarging the receptive field. A larger receptive field is achieved by tuning two hyperparameters: the kernel size (ks) and the dilation factor (dl), where the layer-wise effective history is calculated as $(ks - 1) * dl$. Following convention, we set dl to grow exponentially with depth ($dl = O(2^i)$), a strategy that ensures comprehensive input coverage and facilitates an exceptionally wide effective history in deep networks. An illustrative diagram is provided in Figure 7a.



(**a**) A dilated non-causal convolution. (**b**) Residual block.

Figure 7. Architectural elements in a non-causal TCN. (**a**) A dilated non-causal convolution. Yellow boxes indicate valid features, white boxes indicate padding, the blue box marks the example output, and the blue lines show convolutional connections. (**b**) Residual block. The meaning of the lines is as shown in the legend.

3.3.3. Residual Connections

Within a residual block, a parallel pathway applies a series of transformations F to the input. The result of these transformations is then added back to the block's original input $\mathbf{x}$ via a skip connection.

$$o = \text{Activation}(\mathbf{x} + \mathcal{F}(\mathbf{x})) \tag{16}$$

By structuring learning in this way, each layer is tasked with refining the identity map rather than computing it entirely from scratch. This principle has consistently proven critical for successfully training extremely deep architectures.

The dependence of a TCN's receptive field on its depth nd, filter size ks, and dilation base dl poses a significant challenge: the stabilization of progressively deeper and larger networks must be addressed.

As shown in Figure 7b, TCN uses a residual block containing two dilated non-causal convolution layers coupled with GELU activation. All convolutional filters undergo weight normalization, and further applies spatial dropout for regularization—a procedure that deactivates entire channels randomly at each training step.

Although both architectures employ residual blocks, TCNs diverge from ResNets in their handling of tensor shapes. The direct addition used in ResNets is insufficient when feature map depths differ. Therefore, a 1×1 convolution is introduced as a projection to ensure compatible shapes for the $\oplus$ operation, as depicted in Figure 7b.

4. Experiments and Results

4.1. Orbit Uncertainty Estimation

In the following analysis, we perform phase space embedding assessment and Lyapunov exponent estimation on a 6D ephemeris/state sequence $[x, y, z, v_x, v_y, v_z]$. The se-

quence is obtained by propagating publicly available TLEs using SGP4 and resampling the resulting state at 1 Hz over a 24 h interval (86,400 samples). Unless otherwise stated, AMI/FNN are computed on each state component and then aggregated following the multivariate extension described in Section 2.1.

4.1.1. Experimental Settings

The experiments are conducted on an Ubuntu 22.04 platform equipped with an NVIDIA 4080 GPU (12 GB VRAM). Unlike conventional forecasting tasks strictly constrained by temporal causality, the primary objective of this study is to calibrate low-quality SGP4 orbital data into high-fidelity trajectories. This formulation allows the framework to leverage comprehensive temporal context—incorporating both past and future information from the pre-available SGP4 stream—to refine the state estimate at any given epoch.

Each sample in our dataset is defined as a continuous 24 h ephemeris sequence with a 1 s sampling interval (86,400 steps), ensuring alignment with the typical TLE update cycle. The input and output at each step are represented by a six-dimensional state vector (x, y, z, v_x, v_y, v_z) capturing position and velocity. Due to the practical constraints of TLE data acquisition (source: Space-Track.org), which is typically released only once daily with limited historical accessibility, we compiled a total of 564 samples from 10 satellites over a 20-day period.

High-precision "HPOP truth" trajectories serve as the ground truth, generated via Systems Tool Kit (STK, Ansys, Inc., Canonsburg, PA, USA). (STK 11.6) using a Runge–Kutta fourth-order integrator with a fixed 1 s step. This numerical propagation incorporates complex perturbations such as atmospheric drag (NRLMSISE-00) and non-spherical gravitational effects (JGM-3) to ensure high-fidelity reference states. To ensure experimental rigor and strictly results, we adopted a chronological block-wise partitioning strategy rather than random sample-wise shuffling. The dataset is divided into a training set (80%) and a validation set (20%) based on distinct time intervals, ensuring that the validation orbits are temporally isolated from the training data.

To maintain maximum training efficiency given the dataset size (564 samples derived from 20 days of TLE data), we employed a two-way data split rather than a conventional three-way split. It is important to emphasize that the 20% evaluation subset functions as a strictly independent test set. We did not perform hyperparameter optimization or utilize early stopping mechanisms based on this subset.

Model training is performed for 3000 epochs using the Mean Squared Error (MSE) loss and the AdamW optimizer. A dynamic learning rate strategy is employed, starting at an initial value of 0.001 and gradually decreasing to facilitate finer adjustments and stable convergence.

4.1.2. Performance Evaluation

In Figure 8, we have recorded the training time and convergence rate of different methods. The results indicate that the non-causal DIR-TCN (Non-Causal Dynamics-Informed Representation TCN, Non-Causal DIR-TCN) model outperforms previous models in both training duration and convergence speed. Compared to other methods, the DIR-TCN model we proposed has a faster convergence rate and is more stable after convergence. At the same time, its performance on both the training set and the validation set is superior to other models.

The superior performance illustrated in Figure 8 can be attributed to two key architectural modifications. First, the dynamics-informed encoding within the phase space allows the model to augment the representation of the current state with future information, leading to a more comprehensive perception. This is further refined by fully-connected

layers that prune redundant information, effectively balancing descriptive completeness with computational efficiency. Second, we modified the Temporal Convolutional Network (TCN) by removing its causal constraints, transforming it into a non-causal architecture. This enables the revised TCN to incorporate future contexts for its estimations. The synergy between this non-causal property and the dynamics-informed encoding provides the model with a broader temporal perspective, thereby optimizing the estimation outcomes.

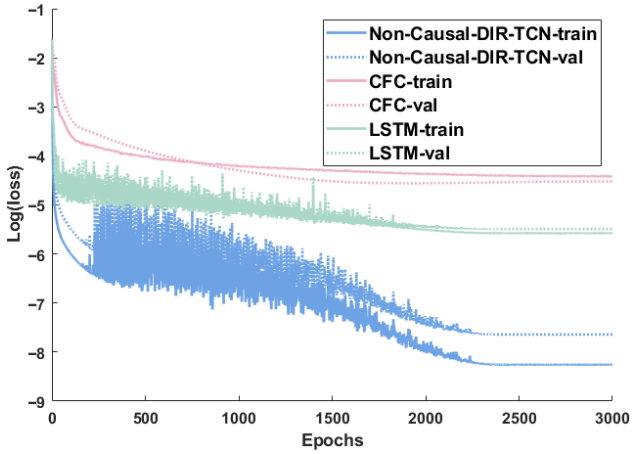


Figure 8. Performance comparison among models.

In contrast, the other methods presented in the figure are confined to strict causal frameworks. They must rely solely on historical data for prediction, inherently limiting their access to future information and consequently constraining their performance.

We denormalized the results and compared them with SGP4 data. However, due to the large magnitude of the orbital values, a direct comparison was not sufficiently informative. Therefore, we evaluated the errors of both methods relative to HPOP data and presented the comparison in Figure 9.

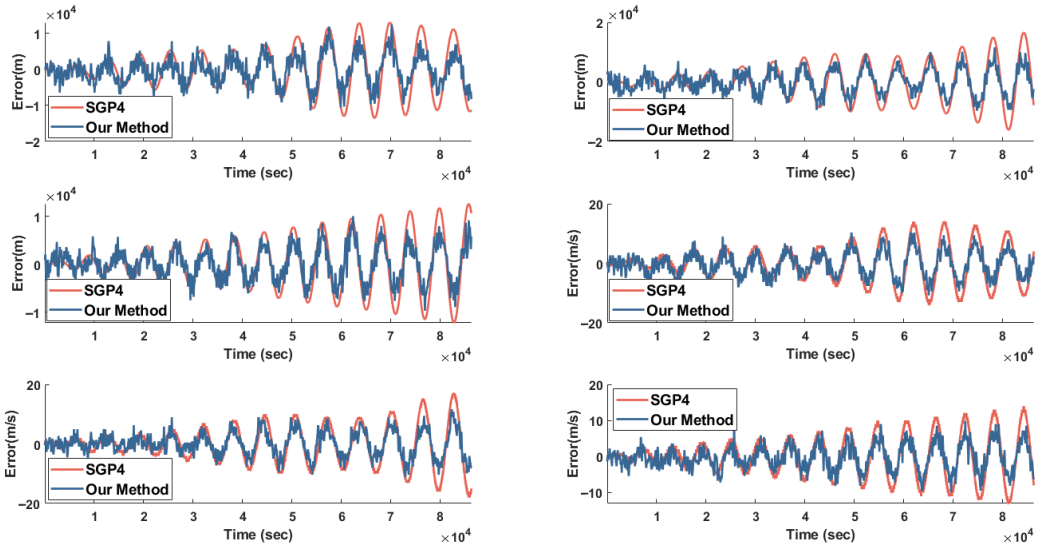


Figure 9. Orbit error comparison for ORBCOMM FM110.

As shown in Figure 9, the error between SGP4 and HPOP primarily accumulates over time, making the discrepancy negligible at the beginning of the propagation. In contrast, our method exhibits considerable fluctuations during the early stage, which might be attributed to the model's need for a short "warm-up" period to establish a stable temporal context. On the other hand, the SGP4 model, being an analytical propagator, is deterministic and

exact at the initial epoch. However, as time progresses, the error accumulation mechanism in SGP4, driven by unmodeled perturbations, dominates over our model's initial transient behavior. To better adapt the approach for opportunistic signal positioning, we integrated both datasets: SGP4 data were used when its error remained small, while our method was employed for trajectory determination when the propagation duration extended and SGP4 accumulated significant errors, as illustrated in Figure 10.

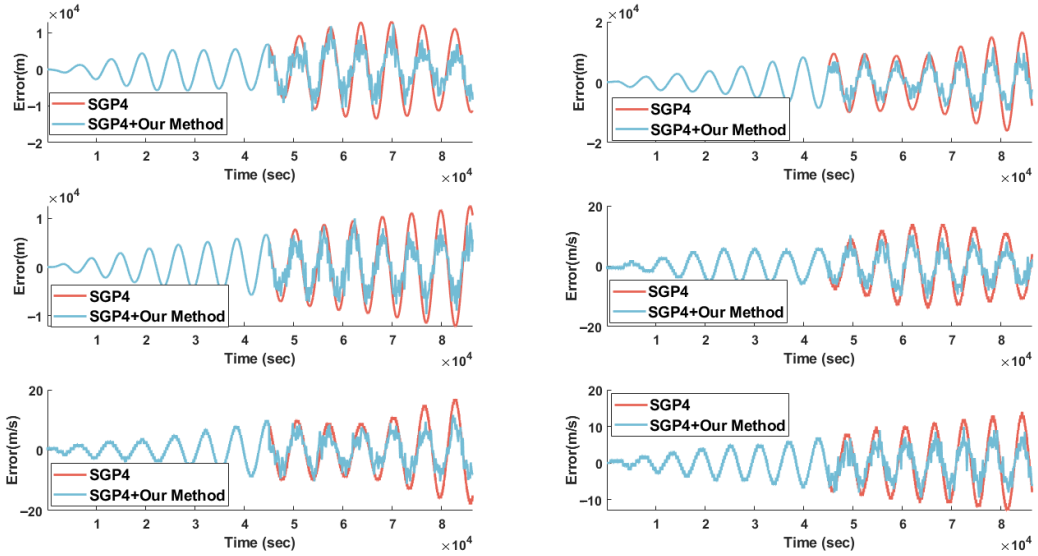


Figure 10. Orbit errors in ORBCOMM FM110 used for SOP positioning.

4.1.3. Ablation Study

In this section, we decouple the Non-Causal DIR-TCN into several modules and perform comprehensive ablation studies. We aim to validate the contribution of each component of Non-Causal DIR-TCN to the extraction of navigation observations.

We evaluated each component of our proposed Non-Causal DIR-TCN under identical experimental settings, and the results are illustrated in Figure 11.

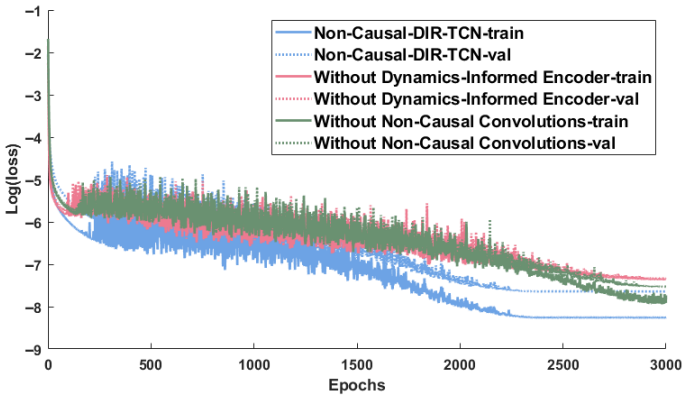


Figure 11. Ablation study results.

Dynamics-Informed Encoder.

As shown in Table 1, we compared the performance of the model with and without physical information encoding. The complete Non-Causal DIR-TCN outperforms the model lacking physical information encoding by 49.23% (on the validation set). The role of physical information encoding is to reconstruct the phase space and provide more comprehensive information to the model while using simple network layers to extract features and avoid redundant input information.

Table 1. Ablation of Non-Causal DIR-TCN.

Model	Params (M)	FLOPs (G)	MSE (1×10^{-8})	
			Train	Valid
Without Dynamics-Informed Encoder	0.10171	34.53278	3.4202	4.6435
Without Non-Causal Convolutions	0.10198	34.62456	5.8447	7.9665
Non-Causal DIR-TCN	0.10198	34.61927	0.5532	2.2855

Non-Causal Convolutions.

In Figure 11, we modified the causal convolutions in the TCN to enable the neural network to incorporate both past and future context simultaneously. As clearly demonstrated in Table 1, this modification to the fusion method resulted in a notable improvement in both the convergence outcome and speed of the model. This enhanced performance can be attributed to the model's improved capacity to capture richer contextual information.

4.2. LEO-Based SOP Positioning Results

The primary objective of our framework design was to enhance the accuracy of positioning results using opportunistic signals. In this section, we discuss how the refined satellite orbit data contribute to improved positioning performance. As illustrated in Figure 12, the initial position was set 100 km away from the true position. Navigation observables were extracted from satellite signals and combined with the optimized satellite orbit data (i.e., the satellite ephemeris) to perform positioning calculations.

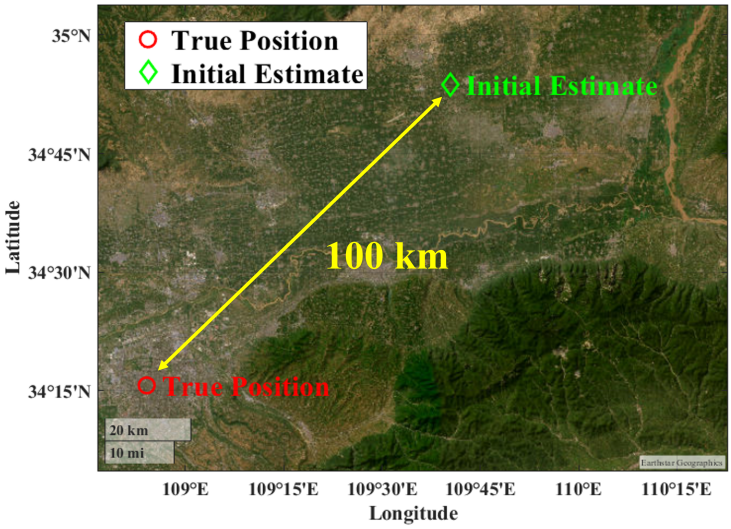


Figure 12. Positioning initialization conditions.

We used HPOP data as the reference orbits and computed the positioning results by inputting the orbits obtained from SGP4 and the non-causal DIR-TCN, respectively. During the positioning process, the Doppler frequency shifts of the satellite signals were employed as navigation observables, with the error set to 3 Hz. The results are presented in Table 2.

Table 2. Error statistic results of SOP positioning with height aid.

	Truth		SGP4		Our Method	
	Mean (m)	RMSE (m)	Mean (m)	RMSE (m)	Mean (m)	RMSE (m)
Eastward error	123.48	153.76	2443.5	2849.1	1968.2	2293.3
Northward error	134.56	168.74	497.40	669.59	402.55	538.20

Table 2 summarizes the positioning errors with elevation assistance (true altitude). When no orbital error is present (i.e., using the reference orbit from HPOP), relatively accurate positioning results can be obtained even in the presence of navigation observables errors, with mean eastward and northward errors both below 150 m. Using the SGP4 orbital data, the mean positioning errors along the x/y-axes are 2443.5 m and 497.40 m, respectively, with significant fluctuations observed in the y-axis (north direction). In comparison, our non-causal DIR-TCN reduces the mean errors to 1968.2 m (x-axis) and 402.55 m (y-axis), demonstrating notable improvement in accuracy along both directions. However, kilometer-level errors persist, primarily due to two factors. First, although our method mitigates certain estimation errors, satellite motion is influenced by multiple forces that cannot be fully modeled using simplified approaches. Moreover, orbital characteristics vary even within the same constellation, limiting generalization to all satellites. Additionally, constraints related to positioning scenarios and hardware restrict model complexity, preventing orbital estimates from fully matching HPOP results and thus affecting positioning accuracy. Second, it can be observed that the eastward error is consistently larger than the northward error regardless of the data used. This disparity is fundamentally caused by the Geometry Dilution of Precision (GDOP) (Appendix A.4) resulting from the orbital design of the satellite constellation. Errors in satellite orbit estimation further exacerbate this inherent discrepancy, making it more pronounced.

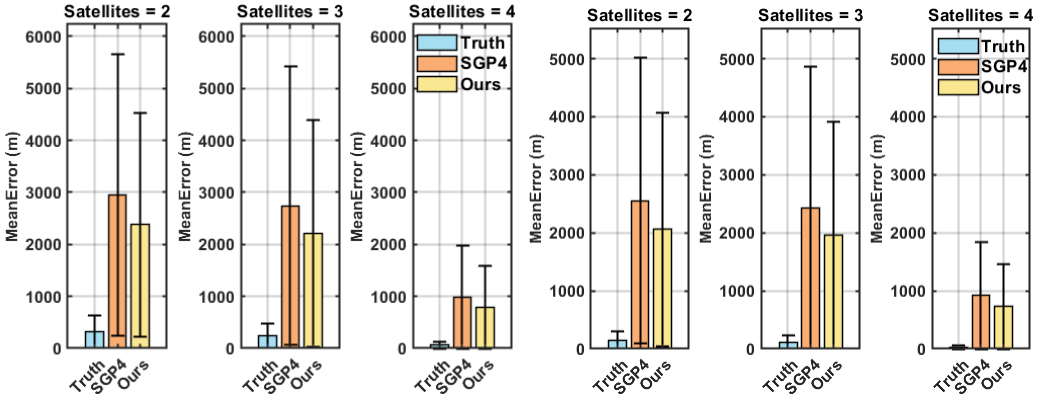
Table 3 confirms that the proposed method outperforms positioning results based on SGP4 orbital data in both 2D and 3D positioning. By reducing errors in satellite orbit estimation, the method enhances the accuracy of both orbit determination and opportunistic signal-based positioning.

Table 3. Comparison of positioning error results with and without height aid.

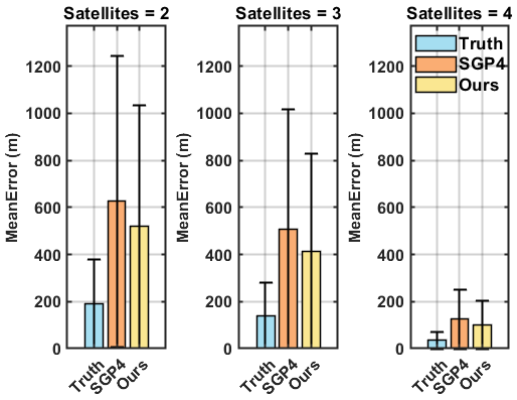
Navigation Observables Methods	3D Mean Error (m)	2D Mean Error (m)
Truth (HPOP)	248.99	202.91
SGP4	2746.8	2570.7
Our Method	2212.4	2070.6

As can be seen from Tables 2 and 3, the method proposed in this paper demonstrates better positioning stability and performance compared to using original SGP4 data under the condition of three-satellite positioning. To further analyze the advantages of the proposed method in orbit estimation and to discuss the impact of satellite geometry (Geometric Dilution of Precision, GDOP) on positioning results, we conducted experiments using 2, 3, and 4 non-cooperative satellites for positioning (the number of simultaneously visible satellites under a single-constellation condition depends on the total number of satellites in the constellation and the orbital design). The experimental results are shown in Figure 13. It can be clearly observed that increasing the number of positioning satellites effectively reduces positioning errors, as the addition of new satellites optimizes the satellite geometry, thereby yielding improved positioning results.

We additionally evaluate the positioning performance in terms of errors relative to the HPOP-derived reference trajectory ("Truth"), rather than only reporting absolute positioning errors. Figure 14 present the relative 3D, eastward, and northward positioning errors under configurations with two, three, and four satellites, respectively.



(**a**) 3D positioning error (**b**) Eastward positioning error



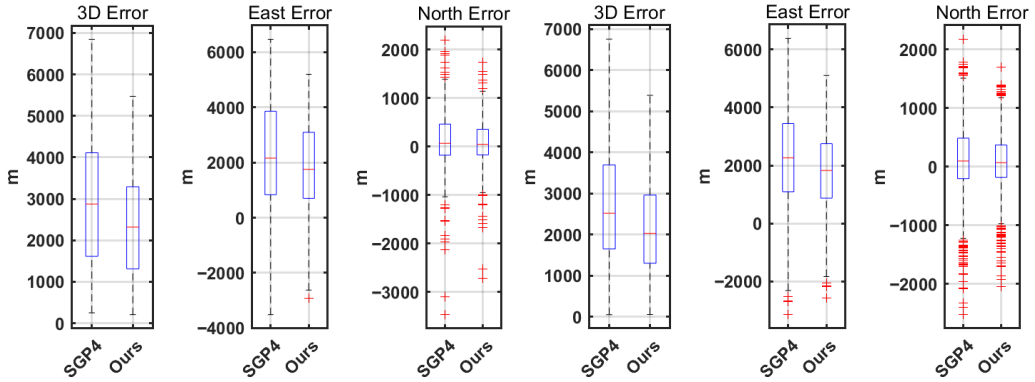
(**c**) Northward positioning error

Figure 13. Comparison of positioning performance with different numbers of satellites. (**a**) 3D positioning error: The 3D error includes the altitude component. When altitude can be obtained by other means (e.g., a barometric altimeter), more accurate positioning results can be achieved. (**b**) Eastward positioning error: This is the primary source of error, which can be effectively reduced by optimizing the satellite geometry. (**c**) Northward positioning error.

Under all satellite configurations, the proposed Non-Causal DIR-TCN consistently achieves smaller median errors and reduced spread compared to SGP4. This demonstrates that refining the orbital states toward HPOP-equivalent trajectories effectively reduces the propagation of ephemeris errors into the positioning solution. The improvement is particularly evident in the eastward component, which is more sensitive to orbital geometry and thus more strongly affected by orbit modeling inaccuracies.

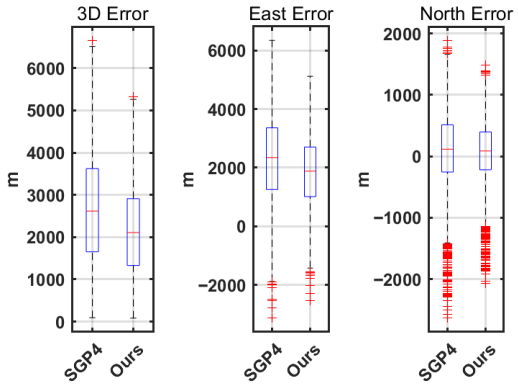
These relative error results further validate that the proposed orbit enhancement framework not only reduces orbital estimation discrepancies but also translates these improvements into measurable gains in SOP-based positioning accuracy.

The primary objective of opportunistic signal-based positioning is originally to address basic position needs in extremely challenging environments—such as GNSS-denied areas or remote regions like deserts and open oceans. However, to achieve higher accuracy, it is often necessary to integrate other positioning technologies, such as inertial navigation (which suffers from error accumulation) and visual navigation (which is susceptible to environmental interference), to further refine the positioning results.



(**a**) 2 satellites positioning error

(**b**) 3 satellites positioning error



(**c**) 4 satellites positioning error

Figure 14. Relative positioning performance with different numbers of satellites evaluated against the HPOP-derived reference trajectory. (**a**) Two-satellite positioning error. (**b**) Three-satellite positioning error. (**c**) Four-satellite positioning error.

5. Conclusions

This paper presents a generalizable deep learning framework designed to enhance low-quality temporal data across a wide range of sensing and measurement applications. Its core innovation lies in a dynamics-informed, non-causal encoding architecture that integrates phase space reconstruction with a modified Temporal Convolutional Network (TCN). This design relaxes temporal causality to capture comprehensive contextual dependencies—a principle directly transferable beyond orbital estimation to any domain involving noisy, low-resolution, or computationally constrained sensor streams. We demonstrate the framework's efficacy by applying it to the specific problem of LEO satellite orbit estimation and subsequent signals of opportunity (SOP) positioning, where it successfully converts low-precision SGP4 data into trajectories of competitive accuracy, serving as a concrete validation of its broader utility for data quality augmentation.

Although our approach does not fully bridge the gap between SGP4 and HPOP data, this limitation stems from several factors: satellite motion is influenced by multiple perturbing forces that cannot be fully captured by simplified models; notable orbital discrepancies exist even among satellites within the same constellation, making complete generalization challenging; and constraints imposed by the positioning scenario and hardware resources limit model complexity. Consequently, our orbit estimation cannot fully replicate HPOP accuracy—in other words, the task requires a model that balances generalization ability and precision under a limited parameter budget. This, in turn, affects final positioning performance.

However, SOP-based positioning is inherently designed to provide basic positioning capabilities in challenging environments (e.g., GNSS-denied areas, deserts, or open oceans) by leveraging non-navigation signals. When higher accuracy is required, our method can be flexibly integrated with other positioning techniques—such as inertial navigation or visual navigation—to achieve robust and refined positioning results.

Author Contributions: Z.W.: writing—original draft, methodology, data curation. H.L.: writing—review and editing, supervision, project administration, data curation, conceptualization. Z.B.: writing—review and editing, validation. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by National Key Laboratory of Unmanned Aerial Vehicle Technology, grant number WRFX202507.

Data Availability Statement: Data will be made available on request.

Conflicts of Interest: The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

Appendix A.1. Doppler Measurement Model

For satellite i, let the satellite position and velocity be $\mathbf{r}_i(t)$ and $\mathbf{v}_i(t)$ from the ephemeris, and the receiver state be $\mathbf{r}(t)$ (and $\mathbf{v}(t)$ if estimated). The line-of-sight unit vector is defined:

$$\mathbf{u}_i(t) = \frac{\mathbf{r}_i(t) - \mathbf{r}(t)}{\|\mathbf{r}_i(t) - \mathbf{r}(t)\|}. \tag{A1}$$

The range rate is

$$\dot{\rho}_i(t) = \mathbf{u}_i^\top(t)\big(\mathbf{v}_i(t) - \mathbf{v}(t)\big), \tag{A2}$$

and the Doppler shift is modeled as

$$f_{D,i}(t) = -\frac{f_c}{c}\,\dot{\rho}_i(t) + \epsilon_i(t), \tag{A3}$$

where f_c is the carrier frequency, c is the speed of light, and $\epsilon_i(t)$ is the Doppler noise (set to 3 Hz in our simulations).

Appendix A.2. Position Estimator

We estimate the horizontal position $\mathbf{r}$ (altitude fixed to the true height in the height-aided setting) by minimizing the Doppler residuals over all observed satellites:

$$\min_{\mathbf{r}} \sum_i \Big(f_{D,i}^{\text{obs}} - f_{D,i}(\mathbf{r})\Big)^2. \tag{A4}$$

We solve it iteratively using Gauss–Newton:

$$\Delta\mathbf{r} = \Big(\mathbf{H}^\top\mathbf{W}\mathbf{H}\Big)^{-1}\mathbf{H}^\top\mathbf{W}\,\mathbf{e}, \tag{A5}$$

where $\mathbf{e}$ stacks the Doppler residuals and $\mathbf{H}$ is the Jacobian matrix with rows $\partial f_{D,i}/\partial\mathbf{r}$.

Appendix A.3. Jacobian

Using $\mathbf{u}_i = \mathbf{s}_i/\rho_i$ with $\mathbf{s}_i = \mathbf{r}_i - \mathbf{r}$ and $\rho_i = \|\mathbf{s}_i\|$, the derivative can be written as

$$\frac{\partial f_{D,i}}{\partial\mathbf{r}} = -\frac{f_c}{c}\,\frac{\partial\dot{\rho}_i}{\partial\mathbf{r}}, \tag{A6}$$

with

$$\frac{\partial \dot{\rho}_i}{\partial \mathbf{r}} = -\left(\frac{\mathbf{I} - \mathbf{u}_i \mathbf{u}_i^\top}{\rho_i}\right)(\mathbf{v}_i - \mathbf{v}). \tag{A7}$$

(For the height-aided static receiver, $\mathbf{v} = \mathbf{0}$ and only the horizontal components of $\mathbf{r}$ are updated).

Appendix A.4. GDOP Definition

We define the geometry dilution of precision as

$$\mathrm{GDOP} = \sqrt{\mathrm{trace}\left(\left(\mathbf{H}^\top \mathbf{W} \mathbf{H}\right)^{-1}\right)}, \tag{A8}$$

which quantifies the sensitivity of the position solution to measurement noise under the given satellite geometry.

References

1. Wu, N.; Qin, H.; Zhao, C. Long-Baseline Differential Doppler Positioning Using Space-Based SOP Based on BPVGMM. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 8503610. [CrossRef]
2. Shahcheraghi, S.; Kassas, Z.M. A Computationally Efficient Approach for Acquisition and Doppler Tracking for PNT with LEO Megaconstellations. *IEEE Signal Process. Lett.* **2024**, *31*, 2400–2404. [CrossRef]
3. Sabbagh, R.; Kassas, Z.M. Observability Analysis of Receiver Localization via Pseudorange Measurements From a Single LEO Satellite. *IEEE Control Syst. Lett.* **2023**, *7*, 571–576. [CrossRef]
4. Farhangian, F.; Landry, R. High-order pseudorange rate measurement model for multi-constellation LEO/INS integration: Case of Iridium-NEXT, Orbcomm, and Globalstar. *Proc. Inst. Mech. Eng. Part G-J. Aerosp. Eng.* **2023**, *237*, 925–939. [CrossRef]
5. Kassas, Z.Z.M.; Morales, J.J.; Khalife, J.J. New-Age Satellite-Based Navigation STAN: Simultaneous Tracking and Navigation with LEO Satellite Signals. *Inside GNSS Mag.* **2019**, *14*, 56–65.
6. Hayek, S.; Saroufim, J.; Kassas, Z.M. Ephemeris Error Correction for Tracking Non-Cooperative LEO Satellites with Pseudorange Measurements. In Proceedings of the 2024 IEEE Aerospace Conference, Big Sky, MT, USA, 2–9 March 2024; pp. 1–9.
7. Khairallah, N.; Kassas, Z.M. Ephemeris Tracking and Error Propagation Analysis of LEO Satellites with Application to Opportunistic Navigation. *IEEE Trans. Aerosp. Electron. Syst.* **2024**, *60*, 1242–1259. [CrossRef]
8. Chen, L.; Bai, X.-Z.; Liang, Y.-G.; Li, K.-B. Orbital Prediction Error Propagation of Space Objects. In *Orbital Data Applications for Space Objects: Conjunction Assessment and Situation Analysis*; Springer: Singapore, 2017; pp. 23–75. [CrossRef]
9. Zaheer, M. Kinematic orbit determination of low Earth orbiting satellites, using satellite-to-satellite tracking data and comparison of results with different propagators. *Geod. Geoinform.* **2014**. Available online: https://api.semanticscholar.org/CorpusID:117900201 (accessed on 10 March 2026).
10. Khoroshylov, S.V.; Redka, M.O. Deep learning for spacecraft guidance, navigation, and control. *Space Sci. Technol.* **2021**, *27*, 38–52. [CrossRef]
11. Silvestrini, S.; Lavagna, M. Deep Learning and Artificial Neural Networks for Spacecraft Dynamics, Navigation and Control. *Drones* **2022**, *6*, 270. [CrossRef]
12. Caldas, F.; Soares, C. Machine learning in orbit estimation: A survey. *Acta Astronaut.* **2024**, *220*, 97–107. [CrossRef]
13. Curzi, G.; Modenini, D.; Tortora, P. Two-line-element propagation improvement and uncertainty estimation using recurrent neural networks. *CEAS Space J.* **2022**, *14*, 197–204. [CrossRef]
14. Li, B.; Zhang, Y.; Huang, J.; Sang, J. Improved orbit predictions using two-line elements through error pattern mining and transferring. *Acta Astronaut.* **2021**, *188*, 405–415. [CrossRef]
15. Bae, T.S.; Kim, M.H.; Kim, M.S.; Lim, S.H. Real-time Orbit Prediction based on Deep Learning. In *AGU Fall Meeting Abstracts*; American Geophysical Union (AGU): Washington, DC, USA, 2019; p. G21C-0753.
16. Varey, J.; Ruprecht, J.D.; Tierney, M.; Sullenberger, R. Physics-Informed Neural Networks for Satellite State Estimation. In Proceedings of the 2024 IEEE Aerospace Conference, Big Sky, MT, USA, 2–9 March 2024; pp. 1–8. [CrossRef]
17. Acciarini, G.; Baydin, A.G.; Izzo, D. Closing the gap between SGP4 and high-precision propagation via differentiable programming. *Acta Astronaut.* **2025**, *226*, 694–701. [CrossRef]
18. Zauchner, M.G.; Horsfield, A.; Lischner, J. Accelerating GW calculations through machine-learned dielectric matrices. *npj Comput. Mater.* **2023**, *9*, 184. [CrossRef]
19. Zhang, Y.; Tian, S.; Liao, M.; Zou, W.; Xu, C. A hybrid domain learning framework for unsupervised semantic segmentation. *Neurocomputing* **2023**, *516*, 133–145. [CrossRef]

20. Loshelder, G.E.; Sood, R.; Su, W. Orbit Determination Using Physics-Informed Neural Networks. In Proceedings of the AIAA SCITECH 2025 Forum, Orlando, FL, USA, 6–10 January 2025. [CrossRef]
21. Lee, J.Y.; Yoo, S.J. Spatio Temporal Approach of Transformer-Based Satellite Orbit Determination Model. In Proceedings of the 2025 International Conference on Artificial Intelligence in Information and Communication (ICAIIC), Fukuoka, Japan, 18–21 February 2025; pp. 0729–0733. [CrossRef]
22. Mortlock, T.; Kassas, Z.M. Assessing Machine Learning for LEO Satellite Orbit Determination in Simultaneous Tracking and Navigation. In Proceedings of the 2021 IEEE Aerospace Conference (50100), Big Sky, MT, USA, 6–13 March 2021; pp. 1–8. [CrossRef]
23. Xie, S.; Li, J.; Cai, J. Improving High-Precision BDS-3 Satellite Orbit Prediction Using a Self-Attention-Enhanced Deep Learning Model. *Sensors* **2025**, *25*, 2844. [CrossRef]
24. Wallot, S.; Mønster, D. Calculation of Average Mutual Information (AMI) and False-Nearest Neighbors (FNN) for the Estimation of Embedding Parameters of Multidimensional Time Series in Matlab. *Front. Psychol.* **2018**, *9*, 1679. [CrossRef] [PubMed]
25. Kennel, M.B.; Brown, R.; Abarbanel, H.D.I. Determining embedding dimension for phase-space reconstruction using a geometrical construction. *Phys. Rev. A* **1992**, *45*, 3403–3411. [CrossRef]
26. Fraser, A.M.; Swinney, H.L. Independent coordinates for strange attractors from mutual information. *Phys. Rev. A* **1986**, *33*, 1134–1140. [CrossRef]
27. Rosenstein, M.T.; Collins, J.J.; Luca, C.J.D. A practical method for calculating largest Lyapunov exponents from small data sets. *Phys. D* **1993**, *65*, 117–134. [CrossRef]
28. Mehdizadeh, S. A robust method to estimate the largest Lyapunov exponent of noisy signals: A revision to the Rosenstein's algorithm. *J. Biomech.* **2019**, *85*, 84–91. [CrossRef] [PubMed]
29. Curtis, H.D. Preliminary Orbit Determination. In *Orbital Mechanics for Engineering Students*, 2nd ed.; Butterworth-Heinemann: Waltham, MA, USA, 2010; pp. 255–317.
30. Uphoff, C. Numerical Averaging in Orbit Prediction. *AIAA J.* **1973**, *11*, 1512–1516. [CrossRef]
31. Leng, M.; Tay, W.P.; See, C.M.S.; Razul, S.G.; Win, M.Z. Modified CRLB for Cooperative Geolocation of Two Devices Using Signals of Opportunity. *IEEE Trans. Wirel. Commun.* **2014**, *13*, 3636–3649. [CrossRef]
32. Aristoff, J.; Poore, A. Implicit Runge-Kutta Methods for Orbit Propagation. In Proceedings of the AIAA/AAS Astrodynamics Specialist Conference, Minneapolis, MN, USA, 13–16 August 2012. [CrossRef]
33. Sullivan, J.; Grimberg, S.; D'Amico, S. Comprehensive Survey and Assessment of Spacecraft Relative Motion Dynamics Models. *J. Guid. Control Dyn.* **2017**, *40*, 1–23. [CrossRef]
34. Shuster, S. A Survey and Performance Analysis of Orbit Propagators for LEO, GEO, and Highly Elliptical Orbits. Master's Thesis, Utah State University, Logan, UT, USA, 2017. [CrossRef]
35. Tian, X.; Chen, G.; Blasch, E.; Pham, K.; Bar-Shalom, Y. Comparison of three approximate kinematic models for space object tracking. In Proceedings of the 16th International Conference on Information Fusion, Istanbul, Turkey, 9–12 July 2013; pp. 1005–1012.
36. Luo, Y.Z.; Yang, Z. A review of uncertainty propagation in orbital mechanics. *Prog. Aerosp. Sci.* **2016**, *89*, 23–39. [CrossRef]
37. Lewis, P. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *arXiv* **2021**, arXiv:2005.11401.
38. Sriram, A.; Jun, H.; Satheesh, S.; Coates, A. Cold Fusion: Training Seq2Seq Models Together with Language Models. *arXiv* **2017**, arXiv:1708.06426.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.