

Article

Continuous Low-Thrust Maneuver Parameter Detection of Non-Cooperative Satellites Based on a Diffusion Model

Kun Zhang ¹, Yanping Zhou ², Yunhan He ³ and Yun Xu ^{3,*}

¹ School of Computer Science and Information Engineering, Hefei University of Technology, No. 485 Danxia Road, Hefei 230601, China; zkun361@gmail.com

² School of Mathematics, Hefei University of Technology, No. 485 Danxia Road, Hefei 230601, China; zhouyp@mail.hfut.edu.cn

³ School of Management, Hefei University of Technology, No. 193 Tunxi Road, Hefei 230009, China; heyh@hfut.edu.cn

* Correspondence: xuyunxy@hfut.edu.cn

Abstract

It is challenging to detect continuous low-thrust maneuver parameters of non-cooperative satellites because the signals are weak over limited observation arcs and are readily masked by measurement noise and orbit-determination errors. This paper proposes a conditional diffusion model for detecting and estimating continuous low-thrust maneuver parameters from relative-orbit observations. The method uses relative-orbit observations of the non-cooperative target to construct conditional inputs that incorporate orbital dynamical priors. Single-step differencing and dimensionless processing are then used to strengthen weak maneuver signatures. The conditional diffusion model learns the evolution of maneuver parameters under noisy conditions and estimates three-axis continuous low-thrust acceleration sequences. Based on simulations considering the Gaussian noise of relative positions and velocities, the proposed method achieved 85.2% maneuver detection accuracy, while that of the batch least-squares benchmark method was 67.8%. The proposed method is simulated and verified based on Sentinel-6A. Results show that the continuous low-thrust maneuver can be robustly identified under low signal-to-noise ratios and the temporal parameter evolution can be recovered. The method provides a practical route for analyzing non-cooperative satellite maneuver and supporting on-orbit space situational awareness.

Keywords: non-cooperative satellite; continuous low-thrust maneuver; parameter detection; conditional diffusion model; space situational awareness



Academic Editor: Fabio Celani

Received: 28 May 2026

Revised: 10 July 2026

Accepted: 13 July 2026

Published: 16 July 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 \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The rapid deployment of low Earth orbit (LEO) mega constellations and the growing number of satellites are making near-Earth orbital resources increasingly congested. To maintain constellation configurations and avoid space debris, satellites must perform orbit-maintenance and collision-avoidance maneuvers more frequently, placing higher demands on space situational awareness systems [1–4]. Detecting maneuver parameters is a key component of space situational awareness [5–7]. Reliable identification of weak maneuver signatures in complex space environments has important engineering value [8,9].

Conventional methods for detecting space-object maneuvers mainly rely on statistical approaches [10–12] or dynamical approaches [13,14]. These methods were developed largely for high-orbit satellites and for impulsive maneuvers generated by large-thrust propulsion. Large-thrust impulsive maneuvers usually appear as abrupt velocity changes

and produce clear variations in orbital elements such as the semimajor axis and eccentricity. They can be detected with preset thresholds. Modern LEO constellation satellites, however, increasingly use electric-propulsion continuous low thrust [15]. Over short arcs, such a maneuver produces only small accumulated orbital deviations [16]. In real observations, these deviations are easily obscured by measurement noise, orbital determination errors and unmodeled perturbations. Distinguishing the weak physical signatures of continuous low-thrust maneuvers from strong background noise and complex errors remains an unresolved challenge in maneuver parameter detection.

Recent advances in deep learning have introduced data-driven temporal networks, including Long Short-Term Memory networks (LSTM) and Convolutional Neural Networks (CNNs), into space-object maneuver detection [17–20]. Existing intelligent detection methods still have two limitations for continuous low-thrust maneuvers. First, data-driven models generally require large volumes of on-orbit data of satellites. Even when simulations are used, most available datasets contain only position, velocity and orbital elements, which limits training efficiency and detection accuracy [21,22]. Second, the acceleration signatures of continuous low thrust are intrinsically weak and are readily masked by measurement noise and model errors. Temporal networks represented by LSTM models struggle to learn these weak maneuver features sufficiently.

To address these limitations, we propose a continuous low-thrust maneuver parameter detection method based on a conditional diffusion model. Conditional diffusion models can learn noise-contaminated sequence distributions and progressively suppress error-induced fluctuations during reverse denoising. To improve data representation and physical constraints, we construct simulated relative-orbit data for model training and use processed Sentinel-6A on-orbit data only for case-level testing. We also expand the feature representation through difference-based normalization and embed relative orbital motion information into the model input. The proposed method enables stable detection and parameter estimation of continuous low-thrust maneuvers over short observation arcs.

This study makes three main contributions. Firstly, single-step differencing and dimensionless processing are applied to relative orbital data to introduce orbital dynamical priors as conditional inputs to the diffusion model. Secondly, a conditional diffusion-based method is developed for continuous low-thrust detection. It suppresses noise by learning parameter distributions under error-contaminated conditions and strengthens weak maneuver signatures. Thirdly, a unified conditional diffusion detection framework is constructed to accurately detect the weak signals and acceleration components that differ substantially in magnitude. This improves the stability and adaptability of maneuver-parameter identification in complex noisy environments.

The remainder of this paper is organized as follows. Section 2 describes the construction of the relative orbital dataset, including the assumptions, dataset generation and dimensionless processing of relative orbital data. Section 3 presents the conditional diffusion model, including the noising and denoising of relative orbital data, the network architecture and the evaluation metrics. Section 4 reports simulation results, a comparison with a batch least-squares benchmark method, ablation experiments and a Sentinel-6A real-data case study. Section 5 concludes the paper.

2. Relative Orbital Data Preparation

2.1. Assumptions and Scenario Definition

To describe absolute and relative spacecraft motion, two coordinate frames are defined. The Earth-centered inertial (ECI) frame has its origin at the Earth's center of mass. Its X-axis points to the vernal equinox, its Z-axis follows the Earth's rotation axis towards the North Pole, and its Y-axis completes a right-handed frame. The position vector of the

target satellite in the ECI frame is expressed as $\mathbf{R} = [X, Y, Z]^T$, where X , Y and Z are the inertial X-, Y- and Z-position components in the ECI frame, respectively. The corresponding velocity is denoted by $\mathbf{V} = [V_X, V_Y, V_Z]^T$, where V_X , V_Y and V_Z are the inertial X-, Y- and Z-velocity components.

The chief-centered Local Vertical Local Horizontal (LVLH) frame has its origin at the chief satellite's center of mass. Its x -axis points radially from the Earth's center to the chief satellite, its z -axis is normal to the orbital plane and aligned with the orbital angular momentum, and its y -axis completes the right-handed triad in the along-track direction, as shown in Figure 1. The relative position of the target with respect to the chief is denoted by $\mathbf{r} = [x, y, z]^T$, where x , y and z are the radial, along-track and normal position components in the LVLH frame, respectively. The corresponding relative velocity is denoted by $\mathbf{v} = [v_x, v_y, v_z]^T$, where v_x , v_y and v_z are the radial, along-track and normal velocity components.




Figure 1. Earth-centered inertial frame and chief-centered LVLH frame.

Before studying maneuver detection of non-cooperative space targets, the assumptions of the observation scenario must be specified. The scenario contains a chief satellite and a non-cooperative target satellite. The chief satellite operates in a near-circular orbit within the typical altitude range of LEO constellation satellites. The non-cooperative target satellite is assumed to use an electric-propulsion system. From the perspective of the chief satellite, the start time, end time, duration and thrust vector of the target maneuver are unknown. During a single spatial maneuver arc, the target thrust-acceleration vector is assumed to remain approximately constant, which is consistent with the quasi-steady acceleration produced by many low-thrust electric-propulsion systems [23].

For relative observation and state acquisition, the chief satellite is assumed to carry sensing payloads such as an optical camera or spaceborne radar, enabling relative navigation and localization of the non-cooperative target. For a typical range-and-angle relative-navigation measurement, the measurement vector obtained by the chief in the LVLH frame is denoted by $\mathbf{b}_k = [\rho_k, \theta_k, \phi_k]^T$, where ρ_k , θ_k and ϕ_k are the relative range, azimuth and elevation at time t_k , respectively. The measurement model is written as Equation (1):

$$\mathbf{b}_k = \begin{bmatrix} \rho_k \\ \theta_k \\ \phi_k \end{bmatrix} = \begin{bmatrix} \sqrt{x^2 + y^2 + z^2} \\ \arctan(y/x) \\ \arcsin(z/\sqrt{x^2 + y^2 + z^2}) \end{bmatrix} + \mathbf{w}_k \quad (1)$$

where $\mathbf{w}_k$ denotes the sensor and the random measurement noise.

After processing with relative navigation algorithms such as the extended Kalman filter, the chief satellite can obtain relatively accurate relative position and velocity data for

the target satellite. The estimated relative positions and velocities serve as the raw inputs for training in this paper.

Although these filtering algorithms suppress part of the initial observation noise, the final relative-state sequence inevitably retains estimation errors caused by sensor measurement limits and navigation-filtering uncertainty. The Gaussian assumption in this study is applied only to these residual relative-position and relative-velocity estimation errors after relative navigation filtering. The noisy relative position and velocity are then used to construct the single-step differences, acceleration-related features and dimensionless conditional inputs.

$$\begin{aligned}\hat{\mathbf{r}}_k &= \mathbf{r}_k + \varepsilon_{r,k}, \\ \hat{\mathbf{v}}_k &= \mathbf{v}_k + \varepsilon_{v,k},\end{aligned}\quad (2)$$

where $\varepsilon_{r,k}$ and $\varepsilon_{v,k}$ denote the integrated estimation error terms in relative position and relative velocity, respectively.

2.2. Construction of the Relative Orbital Dataset

The dataset is constructed based on the relative-state estimate sequence obtained after the navigation-filtering algorithm described in Section 2.1, namely $\hat{\mathbf{r}}_k$ and $\hat{\mathbf{v}}_k$. In addition, the dataset contains the orbital elements of the target satellite at the same epoch, $\mathbf{O}_k = [a_k, e_k, i_k, \Omega_k, \omega_k, M_k]^T$, where a_k , e_k , i_k , Ω_k , ω_k and M_k denote the semimajor axis, eccentricity, inclination, right ascension of the ascending node, argument of perigee and mean anomaly of the target satellite, respectively.

To extract weak continuous-maneuver features from finite observation sequences and introduce orbital dynamical constraints, the single-step difference features of the relative position, relative velocity and orbital elements are applied for training. Relative position and velocity describe the observed state variation, but they do not explicitly encode the physical source of orbital motion. Because spacecraft motion in near-Earth space is governed mainly by central gravity and orbital perturbations, an acceleration-related term is introduced to embed dynamical information into the conditional dataset.

Neglecting the effects of orbital perturbations such as solar pressure, solar gravity, and lunar gravity on LEO satellites, the acceleration used to characterize orbital dynamics is represented by the sum of the central gravitational acceleration $\mathbf{a}_g$, J_2 perturbation $\mathbf{a}_{J_2}$ and atmospheric drag $\mathbf{a}_{drag}$, as shown in Equations (3)–(7).

$$R = |\mathbf{R}| = \sqrt{X^2 + Y^2 + Z^2} \quad (3)$$

$$\mathbf{a}_g = -\frac{G(M+m)}{R^3}\mathbf{R} = -\frac{\mu}{R^3} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \quad (4)$$

$$\mathbf{a}_{J_2} = \frac{3}{2} \frac{J_2 \mu R_E^2}{R^4} \begin{bmatrix} -\sin \Omega \cos i \sin u + \cos \Omega \cos u & \cos \Omega \cos i \sin u + \sin \Omega \cos u & \sin i \sin u \\ -\sin \Omega \cos i \cos u - \cos \Omega \sin u & \cos \Omega \cos i \cos u - \sin \Omega \sin u & \sin i \cos u \\ \sin \Omega \sin i & -\cos \Omega \sin i & \cos i \end{bmatrix} \begin{Bmatrix} \frac{X}{R} \left(5 \frac{Z^2}{R^2} - 1 \right) \\ \frac{Y}{R} \left(5 \frac{Z^2}{R^2} - 1 \right) \\ \frac{Z}{R} \left(5 \frac{Z^2}{R^2} - 3 \right) \end{Bmatrix} \quad (5)$$

$$\mathbf{a}_{drag} = -\frac{1}{2} \rho(h) C_D \frac{A}{m} \|\mathbf{v}_{rel}\| \mathbf{v}_{rel}, \quad (6)$$

$$\mathbf{a} = \mathbf{a}_g + \mathbf{a}_{J_2} + \mathbf{a}_{drag} \quad (7)$$

where R denotes the orbital radius of the target satellite and R_E is the Earth's equatorial radius. $\rho(h)$ is the atmospheric density at orbital altitude h . C_D is the drag coefficient. A

is the effective cross-sectional area. m is the spacecraft mass. $\mathbf{v}_{\text{rel}}$ is the spacecraft velocity relative to the rotating atmosphere. $\mathbf{v}$ and $\mathbf{r}$ are the inertial velocity and position vectors of the spacecraft. $\boldsymbol{\omega}_E$ is the Earth's rotation angular-velocity vector.

In this study, the observed trajectory and the non-maneuver trajectory are defined. The observed trajectory represents the actual orbital state for the current time, while the non-maneuver trajectory is obtained by propagating the observed state of the previous epoch under the non-maneuver assumption. During the orbital propagation, the two-body gravity, J_2 perturbation and atmospheric drag are considered. The solar-radiation pressure and third-body gravity are not explicitly activated because their acceleration magnitudes are dramatically smaller than the low-thrust acceleration. The acceleration evaluated on the observed trajectory is denoted as $\mathbf{a}_{ob}$, and the acceleration evaluated on the non-maneuver trajectory is denoted as $\mathbf{a}_{id}$.

The acceleration-difference feature is computed by directly subtracting the ideal acceleration from the observed acceleration, as shown in Equation (8):

$$\Delta \mathbf{a} = \mathbf{a}_{ob} - \mathbf{a}_{id} \quad (8)$$

The acceleration-difference feature compares the acceleration of the observed trajectory with that of the non-maneuver trajectory. If the target does not maneuver, $\Delta \mathbf{a}$ is mainly caused by numerical propagation or estimation errors and remaining unmodeled perturbations, such as solar radiation pressure and third-body residuals, whose typical order in LEO is about $10^{-7} \sim 10^{-6} \text{ m/s}^2$ under common spacecraft area-to-mass ratios. If it maneuvers, $\Delta \mathbf{a}$ additionally contains the maneuver acceleration. In the present low-thrust simulations, the dominant thrust component is mainly at the $10^{-3} \sim 10^{-5} \text{ m/s}^2$ order, which is larger than the residual perturbation level. This order-of-magnitude contrast helps distinguish maneuver and non-maneuver trajectories through the acceleration-difference feature.

The simulation dataset is used for model training and statistical testing. The real orbital data are derived from publicly available precise ephemeris or orbit-determination products and are used only for the Sentinel-6A case-level verification in Section 4.4. These real data are converted into relative position and velocity information using the corresponding processing procedure.

2.3. Dimensionless Processing of Relative Orbital Dataset

The relative position, relative velocity, acceleration difference and orbital elements have significantly different numerical scales. In a deep-learning framework, such dimensional differences may lead to gradient vanishing or gradient explosion during back-propagation. They may also affect the numerical stability of the forward noising and reverse generation processes of the diffusion model.

To reduce the large differences among data magnitudes and unify the physical dimensions, this paper applies dimensionless processing to the above dataset.

The dimensional normalization of relative motion is defined in Equations (9)–(12):

$$h_c = \sqrt{\mu a_c (1 - e_c^2)}, R_c = \frac{a_c (1 - e_c^2)}{1 + e_c \cos f_c} \quad (9)$$

$$\bar{x} = x/R_c, \bar{y} = y/R_c, \bar{z} = z/R_c \quad (10)$$

$$\bar{v}_x = v_x R_c / h_c, \bar{v}_y = v_y R_c / h_c, \bar{v}_z = v_z R_c / h_c \quad (11)$$

$$\bar{a}_x = a_x R_c^3 / h_c^2, \bar{a}_y = a_y R_c^3 / h_c^2, \bar{a}_z = a_z R_c^3 / h_c^2 \quad (12)$$

where h_c denotes the angular momentum of the chief satellite, and R_c denotes the orbital radius of the chief satellite. The vectors $\bar{\mathbf{r}} = [\bar{x}, \bar{y}, \bar{z}]^T$, $\bar{\mathbf{v}} = [\bar{v}_x, \bar{v}_y, \bar{v}_z]^T$ and $\bar{\mathbf{a}} = [\bar{a}_x, \bar{a}_y, \bar{a}_z]^T$

represent the normalized relative position, relative velocity and relative acceleration, respectively. After normalization and coordinate system transformation, the final dataset can be expressed as $\begin{bmatrix} \bar{\Delta D}, \bar{\Delta a}, \bar{\Delta O} \end{bmatrix}$, and $\bar{\Delta D} = [\bar{\Delta r}, \bar{\Delta v}]$, $\bar{\Delta r}, \bar{\Delta v}$ are the dimensionless relative position and velocity difference.

3. Diffusion Model

3.1. Noising and Denoising of Maneuver-Parameter Sequences

A diffusion model learns a data distribution by defining two coupled processes. The forward process gradually adds Gaussian noise to a clean sample until it approaches a standard normal distribution, whereas the reverse process learns to remove this noise step by step. In this study, the clean sample x_0 denotes the three-axis continuous low-thrust acceleration sequence within an observation window.

Let β_t be the noise-schedule parameter, $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. The forward noising process can be written compactly as

$$q(x_t | x_0) = \mathcal{N}(\sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t)I) \quad (13)$$

This equation means that any noisy state x_t can be sampled directly from the original maneuver-acceleration sequence x_0 by injecting Gaussian noise at the noise level specified by t . The reverse process uses a neural network to predict the noise component ε contained in x_t under the constraint of relative-orbit observations.

The training objective is the standard noise-prediction loss:

$$L = \mathbb{E}_{x_0, t, \varepsilon} [\|\varepsilon - \varepsilon_\theta(x_t, t, C)\|^2], \quad \varepsilon \sim \mathcal{N}(0, I) \quad (14)$$

where $C = \begin{bmatrix} \bar{\Delta D}, \bar{\Delta a}, \bar{\Delta O} \end{bmatrix}$ is the conditional input sequence constructed in Section 2 and ε_θ is the noise predicted by the network. By minimizing this loss, the model learns how maneuver-parameter sequences are degraded by noise and how they can be recovered from noisy observations.

3.2. Design of the Diffusion Model Network

The proposed network is a conditional diffusion framework for continuous low-thrust maneuver parameter detection. Its role is to estimate the clean maneuver-acceleration sequence x_0 from a noisy diffusion state x_t while using the relative-orbit feature sequence C as a physical condition. The conditional input at epoch k is composed of the dimensionless features $\bar{\Delta r}_k, \bar{\Delta v}_k$ and $\bar{\Delta a}_k$ together with the processed orbital-element vector $\bar{\Delta O}_k$.

The overall framework is shown in Figure 2. The time-series encoder maps the conditional input sequence into a high-dimensional embedding H_c . During reverse denoising, the decoder receives x_t , the diffusion step t and the embedding H_c , and predicts the injected noise:

$$\hat{\varepsilon} = \varepsilon_\theta(x_t, t, H_c) \quad (15)$$

The encoder structure is shown in Figure 3. The normalized multivariate sequence is first transformed by a fully connected layer and then processed by cascaded dilated-convolution blocks. Dilated convolutions enlarge the temporal receptive field without substantially increasing the number of parameters, which helps the model capture long-range relative-orbit evolution. Residual connections are used inside the blocks to stabilize deep feature extraction.




Figure 2. Framework diagram of maneuver parameter detection.




Figure 3. The encoder structure.

After the cascaded blocks, a convolutional layer compresses the features into the conditional embedding H_C . An auxiliary coarse-output branch is retained at the encoder output and supervised during training so that early-stage physical features in the relative-orbit sequence are more directly constrained.

The decoder structure is shown in Figure 4. At each reverse step, the noisy acceleration sequence x_t is aligned by an initial convolutional layer, and the timestep t is encoded using sinusoidal positional embedding to indicate the current noise level. The decoder then fuses the noisy state with the conditional embedding H_C through cross-attention blocks.




Figure 4. The decoder structure.

The denoising update can be viewed as using the predicted noise $\hat{\epsilon}$ to move from x_t toward a cleaner state:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \alpha_t}} \epsilon_{\theta}(x_t, t, z) \right) + \sigma_t \eta, \quad \eta \sim \mathcal{N}(0, I) \quad (16)$$

Here $\sigma_t \eta$ is the sampling noise term and is omitted at the final step. This formulation makes the role of the framework explicit: the encoder supplies the physical relative-orbit condition, and the decoder uses that condition to guide each denoising step toward a physically plausible maneuver-acceleration sequence.

Inside each cross-attention block, the noisy feature sequence provides the query, while the encoder embedding H_c provides the keys and values. The attention weights therefore determine which relative-orbit features are most informative for the current denoising step. This design allows weak thrust signatures that are hidden in noisy acceleration residuals to be repeatedly reinforced during the reverse diffusion process.

The final convolutional layer maps the denoised features to the estimated target sample $\hat{x}_0$, which represents the radial, along-track and normal continuous low-thrust acceleration sequences. To improve training stability under low signal-to-noise ratios, the model uses an exponentially decaying learning-rate schedule and gradient clipping. Because the three acceleration axes have different magnitudes and gradient characteristics, staged axis-wise training is used, and the optimizer is rebuilt when switching axes. A boundary-alignment post-processing step then extracts the common active interval from the three-axis predictions to reduce spurious edge detections and improve maneuver-duration consistency.

3.3. Evaluation Metrics

The continuous acceleration sequence predicted by the diffusion model is further processed by threshold extraction and integration to assess detection accuracy, timeliness and parameter-regression accuracy for engineering use. To avoid redundancy and objectively measure model performance, we use four core metrics: maneuver detection accuracy (MDA), maneuver detection delay, mean velocity-increment deviation and false-alarm rate.

Maneuver detection accuracy (MDA):

$$Acc = \frac{N_{det}}{N_{true}} \quad (17)$$

where N_{det} is the number of successfully detected maneuvers and N_{true} is the total number of true maneuver events.

Mean maneuver detection delay:

$$\Delta T = \frac{1}{N_{det}} \sum_{i=1}^{N_{det}} |T_i^{pred} - T_i^{true}| \quad (18)$$

where T_i^{pred} is the detected occurrence time of the i th maneuver and T_i^{true} is the true occurrence time of the i th maneuver. If a true maneuver is not detected, it is excluded from the delay calculation.

Mean velocity-increment deviation:

$$\Delta V = \frac{1}{N_{det}} \sum_{i=1}^{N_{det}} |\Delta v_i^{pred} - \Delta v_i^{true}| \quad (19)$$

where ΔV is the mean velocity-increment deviation, defined as the difference between the mean detected velocity increment and the true velocity increment. The term Δv_i^{pred} is the detected velocity increment of the i th maneuver, and Δv_i^{true} is the true velocity increment of the i th maneuver. If a true maneuver is not detected, it is excluded from the velocity-increment deviation calculation.

False-alarm rate:

$$FAR = \frac{N_{false}}{N_{pred}} \quad (20)$$

where FAR is the mean false-alarm rate, N_{false} is the number of non-maneuver intervals incorrectly classified as maneuvers, and N_{pred} is the total number of true non-maneuver intervals.

4. Simulation Results

4.1. Simulation of Continuous Low-Thrust Maneuver Detection

To evaluate the proposed conditional diffusion model, simulated continuous low-thrust maneuver arcs were generated and tested under the same observation assumptions as those used for model training. The test set contains maneuver and non-maneuver arcs, allowing both maneuver detection and three-axis acceleration regression to be evaluated.

In Sections 4.1–4.3, the noisy condition refers to zero-mean Gaussian state-estimation errors added to the relative position and relative velocity before single-step differencing, acceleration-feature construction and dimensionless processing. The standard deviation was 1 km for each relative-position axis and 0.1 m/s for each relative-velocity axis in the radial, along-track and normal directions. The maneuver-acceleration labels were not perturbed and were used only as training or evaluation targets.

Figure 5 shows a maneuver-parameter estimation result. The predicted curves recover the maneuver window and the main three-axis acceleration levels, with the y-track component being the dominant thrust direction and the x- and z-track components being much weaker. The weak components are more sensitive to state-estimation noise and residual modeling errors.

Some predicted profiles exhibit small zig-zag fluctuations. The estimates still oscillate around an approximately constant thrust level within the maneuver arc. Fluctuations are mainly attributed to sample-wise prediction residuals and to the forward-noising and reverse-denoising process of the diffusion model. Because adjacent epochs are learned as a temporally correlated sequence, an over- or underestimation at one epoch can influence the local correction at neighboring epochs, producing small compensating oscillations around the constant thrust level.




Figure 5. Illustration of maneuver-parameter detection results.

Figure 6 summarizes the statistical performance over 200 test arcs, including classification status, mean acceleration components, maneuver-start delay and velocity-increment error.




Figure 6. Statistical analysis of detection results.

As shown in Figure 6a, the model detected 115 true maneuver arcs, missed 20 maneuver arcs, produced no false alarms and correctly rejected 65 non-maneuver arcs. The corresponding detection accuracy was 85.2%, and the false-alarm rate was 0%.

Figure 6b shows that the predicted mean acceleration components follow the same magnitude pattern as the ground truth: the y-track component is dominant, whereas the x- and z-track components are weaker. The remaining deviations are more visible in


the weak directions because their true signal levels are close to the noise and residual-error background.

Figure 6c indicates that the maneuver-start error was concentrated near zero, with a mean delay of +0.3 steps, a median of 0 and a standard deviation of about 2.0 steps. Figure 6d further shows that the velocity-increment error remained limited for the detected maneuvers, although the dominant y-track component naturally contributed the largest absolute error.


4.2. Comparison with a Batch Least-Squares Processing Method

At present, there is no generally accepted standard benchmark specifically for the detection of continuous low-thrust maneuver parameters in non-cooperative satellites. To provide a reproducible physics-based comparison, we used the orbital maneuver characteristic-parameter detection method of Deng et al. [5] as a batch least-squares benchmark. This method was originally developed for maneuver characteristic-parameter estimation from short-arc historical relative-orbit data.

As summarized in Figure 7, the benchmark first checks whether the local orbital curvature is anomalous and then defines a moving maneuver-detection window. Historical states before and after the reference interval are propagated forward and backward, respectively, and the two propagated state sequences are fused by weighted batch least squares. The residual between the fused forward and backward sequences is then used to determine the maneuver interval and to estimate the start time, end time and thrust acceleration.



(a) Overall maneuver characteristic-parameter detection.



(b) Bidirectional propagation and weighted batch least-squares fusion.

Figure 7. Flowchart of the batch least-squares benchmark method.

Three comparison groups were evaluated. Benchmark Clean applies the batch least-squares benchmark to noise-free relative-state inputs. Benchmark Noisy applies the same benchmark after injecting the relative-state errors defined in Section 4.1. Diffusion Noisy applies the proposed conditional diffusion model under the same noisy conditions. The results are shown in Figures 8–10 and summarized in Table 1.




Figure 8. Batch least-squares benchmark results without state-estimation noise.




Figure 9. Batch least-squares benchmark results with state-estimation noise.




Figure 10. Comparison between the batch least-squares benchmark and diffusion methods.

Table 1. Confusion matrix.

Metric	Benchmark Clean	Benchmark Noisy	Diffusion Noisy
TP	128	40	115
FN	4	19	20
FP	3	125	0
TN	65	16	65
Detection accuracy	97.0%	67.8%	85.2%
FAR	4.4%	88.7%	0%

The four entries are defined at the test-arc level using interval overlap. TP (true positive) denotes a true maneuver arc in which the detected maneuver interval overlaps the reference maneuver interval. FN (false negative) denotes a true maneuver arc where no maneuver interval is detected. FP (false positive) denotes a detected maneuver interval that cannot be matched to a reference maneuver interval, including detections in non-maneuver arcs and temporally misaligned detections in maneuver arcs. TN (true negative) denotes a non-maneuver arc for which no maneuver interval is detected.

Under the ideal noise-free condition, the batch least-squares benchmark detected most maneuver arcs and recovered the main maneuver interval, as shown in Figure 8. Table 1 reports 128 true positives, 4 missed detections and 3 false alarms, corresponding to 97.0% detection accuracy and a 4.4% false-alarm rate. This verifies that the benchmark can work under favorable input conditions.

After relative-state estimation errors were added, the benchmark degraded markedly. As shown in Figure 9 and Table 1, Benchmark Noisy produced only 40 true positives, 19 missed detections and 125 false alarms, reducing the detection accuracy to 67.8% and increasing the false-alarm rate to 88.7%. The estimated acceleration profiles also deviated strongly from the true maneuver interval and component magnitudes.

By contrast, Diffusion Noisy retained much better stability under the same noisy conditions. It produced 115 true positives, 20 missed detections and no false alarms, giving 85.2% detection accuracy and a 0% false-alarm rate. Figure 10 further shows that its start-

time errors were concentrated near the true boundary, whereas the noisy benchmark had a much wider delay distribution.

For parameter recovery, Figure 10b,d show that the diffusion model kept the mean acceleration components and velocity-increment errors closer to the ground truth than Benchmark Noisy. The improvement is most evident under noisy inputs, where the batch least-squares residual is easily dominated by state-estimation errors.

Overall, the batch least-squares benchmark is effective under clean relative-state inputs but is sensitive to input degradation. The proposed diffusion model provides stronger false-alarm control, temporal localization and parameter-recovery stability when maneuver signatures are weak and noisy.

4.3. Ablation of the Acceleration-Difference Feature

To isolate the role of the acceleration-difference feature, an ablation experiment was performed under the same noisy condition. The baseline model used the full conditional input, including single-step differences, dimensionless orbital features and the acceleration-difference feature. The ablation model kept the same diffusion architecture and evaluation procedure but removed the difference-based input construction, directly using the observed relative position, observed relative velocity and observed acceleration as conditional inputs.

The representative result in Figure 11 and the statistics in Figure 12 and Table 2 show that removing the acceleration-difference feature reduced both maneuver detection stability and parameter-recovery quality.

Compared with the baseline model, the ablation model reduced the number of true positives from 115 to 93 and increased the number of missed detections from 20 to 42. The detection metric decreased from 85.2% to 68.9%.

The ablation model also produced nine false alarms, whereas the baseline model produced none. Its false-alarm rate therefore increased from 0% to 13.8%, indicating that the acceleration-difference feature helps separate true maneuver responses from non-maneuver disturbances.




Figure 11. Ablation experiment test results.




Figure 12. Analysis of ablation results.

Table 2. Overall comparison.

Metric	Baseline Model	Ablation Model	Change
TP	115	93	−22 decrease
FN	20	42	+22 increase
FP	0	9	+9 increase (from zero)
TN	65	56	−9 decrease
Precision	85.2%	68.9%	−16.3 pp decrease
FAR	0%	13.8%	+13.8 pp increase
Delay μ	0.3 steps	10.3 steps	+10.0 steps increase
Delay σ	2.0 steps	58.7 steps	+56.7 steps increase

The temporal and regression statistics show the same tendency. The mean start-time delay increased from 0.3 to 10.3 steps, and the delay standard deviation increased from 2.0 to 58.7 steps. The velocity-increment error in Figure 12d was also larger for the ablation model, especially in the dominant maneuver direction.

These results indicate that direct observed states are less effective than acceleration-difference residuals for representing weak continuous-thrust features. The residual construction provides a clearer mismatch between the observed trajectory and the non-maneuver trajectory, thereby improving detection, false-alarm control and temporal localization.

4.4. Verification Based on Sentinel-6A

A Sentinel-6A case study was used to examine the method using processed real-orbit data. This test follows the observation assumption in Section 2.1: relative navigation has already provided relative position and velocity estimates, and the relative-navigation algorithm itself is outside the scope of this paper.

Sentinel-6A operates in a near-circular orbit at an altitude of approximately 1336 km and performs low-thrust orbit-maintenance maneuvers. The case was constructed from public orbit products and maneuver information from EUMETSAT, the Copernicus program

and International DORIS Service products [24–26]. The public orbit information was converted into the chief-centered LVLH relative-state form used by the model.

The reference maneuver started at 05:15:42 and ended at 05:24:27, corresponding approximately to the sampling interval $k = 136$ –145. The diffusion model predicted a maneuver start time of 05:14:00 and an end time of 05:26:00.

The predicted window was slightly wider than the reference maneuver window, with a one-sided error of approximately 1–2 min. This is acceptable for case-level verification because low-thrust orbit-maintenance maneuvers may include gradual thrust build-up and decay, and the relative orbit responds continuously to this process.

Figure 13 and Table 3 summarize the three-axis acceleration estimates for this case.




Figure 13. Three-direction prediction results.

Table 3. Three-direction prediction statistics.

Direction	Predicted Maximum	True Acceleration	Ratio	Confidence
x	$5.94 \times 10^{-5} \text{ m/s}^2$	$-9.65 \times 10^{-7} \text{ m/s}^2$	61	Weak (very small signal)
y	$7.39 \times 10^{-3} \text{ m/s}^2$	$9.84 \times 10^{-3} \text{ m/s}^2$	0.91	High
z	$5.50 \times 10^{-5} \text{ m/s}^2$	$1.79 \times 10^{-5} \text{ m/s}^2$	3.1	Moderate

The y-track direction was the dominant direction of the Sentinel-6A maneuver. In this direction, the predicted amplitude reached approximately 91% of the reference value, and the predicted time window was consistent with the reference maneuver interval. Thus, the model recovered the dominant-direction maneuver trend well.

The x- and z-track reference accelerations were much weaker, and their amplitude ratios were less reliable. Table 3 shows ratios of about 61 in the x direction and 3.1 in the z direction, indicating that weak-component amplitude recovery still contains obvious bias. These weak directions should therefore be interpreted as qualitative responses to maneuver existence rather than precise acceleration estimates.

Nevertheless, the weak-component predictions remained bounded and were temporally aligned with the maneuver window. The Sentinel-6A case therefore supports the

applicability of the method at the processed relative-state level, especially for identifying the dominant along-track low-thrust maneuver.

5. Conclusions

This study proposed a unified conditional diffusion model for continuous low-thrust maneuver-parameter detection of non-cooperative satellites over finite observation arcs. To enhance weak maneuver signatures, relative-orbit dynamical priors were introduced through single-step state differencing, acceleration-difference residual construction and dimensionless processing, providing stable conditional inputs for diffusion-model learning. The acceleration-difference feature compares the observed trajectory with an ideal no-thrust trajectory propagated under central gravity, J_2 perturbation and atmospheric drag, thereby embedding orbital-dynamical information into the dataset. The proposed framework combines a conditional encoder, a cross-attention decoder, staged training and boundary-alignment processing to recover three-axis continuous low-thrust acceleration sequences under noisy relative-orbit observations. Simulation results show that the method can identify continuous low-thrust maneuvers under low signal-to-noise ratios and recover the temporal evolution of maneuver parameters. Compared with the batch least-squares benchmark method, it shows better false-alarm control, temporal-localization ability and parameter-recovery stability under strong noise. Ablation experiments further indicate that the acceleration-difference residual contributes to weak-signal detection. The Sentinel-6A case study based on processed public orbit products shows that the method can identify the reference maneuver time window and recover amplitude changes in the dominant along-track direction, although biases remain for extremely weak x and z components. The present study assumes an approximately constant thrust-acceleration vector within a single maneuver arc and uses processed relative-state inputs. Future work will consider more complex time-varying thrust, multiple continuous maneuvers, richer raw-observation conditions and lightweight deployment for real-time engineering applications.

Author Contributions: Conceptualization, K.Z., Y.Z. and Y.X.; methodology, K.Z., Y.Z. and Y.X.; software, K.Z.; validation, K.Z., Y.X. and Y.H.; formal analysis, K.Z., Y.Z. and Y.X.; investigation, K.Z. and Y.Z.; resources, K.Z.; data curation, K.Z.; writing—original draft preparation, K.Z., Y.Z. and Y.X.; writing—review and editing, K.Z., Y.Z. and Y.X.; visualization, K.Z., Y.X. and Y.H.; supervision, Y.X. and Y.H.; project administration, Y.X. and Y.H.; funding acquisition, K.Z., Y.X. and Y.H. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Natural Science Foundation of Anhui Province (No. 2508085MA009) and the Fundamental Research Funds for the Central Universities, China (No. JZ2025HGTA0144).

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Conflicts of Interest: The authors declare no conflicts of interest.

Nomenclature

Symbol	Meaning
$R = [X, Y, Z]^T, V = [V_X, V_Y, V_Z]^T$	The position and velocity vector of the target satellite in the ECI frame
$r = [x, y, z]^T, v = [v_x, v_y, v_z]^T$	The relative position and velocity vector of the target satellite in the LVLH frame
$b_k = [\rho_k, \theta_k, \phi_k]^T$	Relative navigation measurement vector at the k th time instant

w_k	Random measurement noise
$\hat{r}_k, \hat{v}_k$	Estimated relative position and relative velocity after navigation filtering
$\varepsilon_{r,k}, \varepsilon_{v,k}$	Comprehensive estimation error of relative position and relative velocity at the k th time instant
$\hat{r}_k, \hat{v}_k$	Estimated error components of relative position and relative velocity at the k th time step
$O_k = [a_k, e_k, i_k, \Omega_k, \omega_k, M_k]^T$	Orbital element vector of the target satellite at the k th time step
a_g	Two-body gravitational acceleration
a_{J_2}	Perturbation acceleration caused by Earth's J_2 term
a_{drag}	Perturbation acceleration caused by the atmospheric drag
Δa	Acceleration difference feature
$\bar{\Delta r}, \bar{\Delta v}$	The dimensionless relative position and velocity difference
$\bar{r} = [\bar{x}, \bar{y}, \bar{z}]^T, \bar{v} = [\bar{v}_x, \bar{v}_y, \bar{v}_z]^T, \bar{a} = [\bar{a}_x, \bar{a}_y, \bar{a}_z]^T$	The dimensionless vectors of relative position velocity and acceleration
C	Conditional input sequence
x_0	Original maneuver acceleration sequence
x_t	Noisy state at diffusion step t
ε	Injected Gaussian noise
ε_θ	Noise predicted by the neural network
H_c	Conditional embedding output by the encoder

References

1. Kennewell, J.A.; Vo, B.N. An overview of space situational awareness. In *Proceedings of the 16th International Conference on Information Fusion*; IEEE: New York, NY, USA, 2013; pp. 1029–1036.
2. Mark, C.P.; Kamath, S. Review of active space debris removal methods. *Space Policy* **2019**, *47*, 194–206. [[CrossRef](#)]
3. Hakima, H.; Bazzocchi, M.C.; Emami, M.R. A deorbiter CubeSat for active orbital debris removal. *Adv. Space Res.* **2018**, *61*, 2377–2392. [[CrossRef](#)]
4. Kazemi, S.; Azad, N.L.; Scott, K.A.; Oqab, H.B.; Dietrich, G.B. Orbit determination for space situational awareness: A survey. *Acta Astronaut.* **2024**, *222*, 272–295. [[CrossRef](#)]
5. Deng, H.; Zhang, X.; Gong, B.; Han, F.; Ning, X. Orbital Maneuvering Characteristic Parameters Detection Method for Space Non-cooperative Targets. *J. Astronaut.* **2025**, *46*, 168–180. (In Chinese) [[CrossRef](#)]
6. Bobrinsky, N.; Del Monte, L. The space situational awareness program of the European Space Agency. *Cosm. Res.* **2010**, *48*, 392. [[CrossRef](#)]
7. Holzinger, M.J.; Luu, K.K.; Sabol, C.; Hill, K. Uncorrelated-track classification, characterization, and prioritization using admissible regions and Bayesian inference. *J. Guid. Control Dyn.* **2016**, *39*, 2469–2484. [[CrossRef](#)]
8. Sajjad, N.; Mirshams, M.; Hein, A.M. Spaceborne and ground-based sensor collaboration: Advancing resident space objects' orbit determination for space sustainability. *Astrodynamics* **2024**, *8*, 325–347. [[CrossRef](#)]
9. Tasif, T.H.; Hippelheuser, J.E.; Elgohary, T.A. Analytic continuation extended Kalman filter framework for perturbed orbit estimation using a network of space-based observers with angles-only measurements. *Astrodynamics* **2022**, *6*, 161–187. [[CrossRef](#)]
10. Yang, S.; Jin, X.; Gong, B.; Han, F. Space-based passive orbital maneuver detection algorithm for high-altitude situational awareness. *Aerospace* **2024**, *11*, 563. [[CrossRef](#)]
11. Long, X.; Leping, Y.; Weiwei, C.; Jinghong, L. GEO spacecraft maneuver detection based on causal inference. *Adv. Space Res.* **2023**, *72*, 3756–3769. [[CrossRef](#)]
12. Kaderali, S.; Misra, A. Detection and characterisation of unknown maneuvers in spacecraft. *Acta Astronaut.* **2023**, *205*, 310–318. [[CrossRef](#)]
13. Goff, G.M.; Black, J.T.; Beck, J.A. Tracking maneuvering spacecraft with filter-through approaches using interacting multiple models. *Acta Astronaut.* **2015**, *114*, 152–163. [[CrossRef](#)]

14. Kelecý, T.; Jah, M. Detection and orbit determination of a satellite executing low thrust maneuvers. *Acta Astronaut.* **2010**, *66*, 798–809. [[CrossRef](#)]
15. Long, J.; Zhang, T. Pillars of space traffic management in the era of LEO mega-constellations: A global perspective. *Adv. Space Res.* **2024**, *74*, 800–816. [[CrossRef](#)]
16. Maisonobe, L.; Parraud, P. Very low thrust station-keeping for low Earth orbiting satellites. *Adv. Space Res.* **2023**, *71*, 1558–1593. [[CrossRef](#)]
17. Perovich, N.; Folcik, Z.; Jaimes, R. Applications of artificial intelligence methods for satellite maneuver detection and maneuver time estimation. In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference (AMOS), Maui, HI, USA, 27–30 September 2022.
18. Zhou, X.; Qin, T.; Ji, M.; Qiao, D. A LSTM assisted orbit determination algorithm for spacecraft executing continuous maneuver. *Acta Astronaut.* **2023**, *204*, 568–582. [[CrossRef](#)]
19. Ye, Y.; Zhu, X.; Shen, X.; Chen, X.; Li, L.; Qin, S.J. Diffusion-based Method for Satellite Pattern-of-Life Identification. *arXiv* **2024**, arXiv:2412.10814.
20. Temple, D.; Poole, M.; Camp, M. Network Enabled-Unresolved Residual Analysis and Learning (NEURAL). In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference; Maui Economic Development: Kihei, HI, USA, 2017.
21. Shorten, D.P.; Humphries, M.; Maclean, J.; Yang, Y.; Roughan, M. Optimal Proposal Particle Filters for Detecting Anomalies and Manoeuvres from Two Line Element Data. *Acta Astronaut.* **2025**, *228*, 709–723. [[CrossRef](#)]
22. Li, F.; Zhao, Y.; Zhang, J.; Zhang, Z.; Wu, D. A station-keeping maneuver detection method of non-cooperative geosynchronous satellites. *Adv. Space Res.* **2024**, *73*, 160–169. [[CrossRef](#)]
23. Conway, B.A. *Spacecraft Trajectory Optimization*; Cambridge University Press: New York, NY, USA, 2010; pp. 1–35.
24. Donlon, C.J.; Cullen, R.; Giulicchi, L.; Vuilleumier, P.; Francis, C.R.; Kuschnerus, M.; Simpson, W.; Bouridah, A.; Caleno, M.; Bertoni, R.; et al. The Copernicus Sentinel-6 mission: Enhanced continuity of satellite sea level measurements from space. *Remote Sens. Environ.* **2021**, *258*, 112395. [[CrossRef](#)]
25. Montenbruck, O.; Hackel, S.; Wermuth, M.; Zangerl, F. Sentinel-6A precise orbit determination using a combined GPS/Galileo receiver. *J. Geod.* **2021**, *95*, 109. [[CrossRef](#)]
26. Willis, P.; Lemoine, F.G.; Moreaux, G.; Soudarin, L.; Ferrage, P.; Ries, J.; Otten, M.; Saunier, J.; Noll, C.; Biancale, R.; et al. The International DORIS Service (IDS): Recent developments in preparation for ITRF2013 [C]/ /IAG 150 Years. In *International Association of Geodesy Symposia*; Springer: Cham, Switzerland, 2015; pp. 631–640.

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.