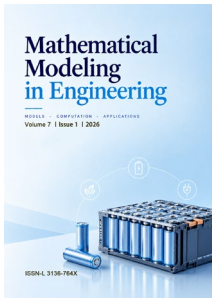
UNIVERSITY OF WEST ATTICA

eISSN: 3136-764X

**Mathematical Modeling in Engineering**

Vol. 7, No. 1, January, 2026




# TimesFM GAT Multisource Fusion Model for Trajectory Correction of Reentry Vehicles

Marek Komorowski[a,*], Sebastian Zalewski[a]

[a] *Faculty of Mechatronics and Automation, Koszalin University of Technology, Koszalin, 75-453, Poland*

Received 29 November 2025; Revised 17 February 2026; Accepted 22 February 2026; Published 27 February 2026

**Abstract**

Trajectory correction during the re-entry phase has problems due to changes in the noise level and time resolution of inertial measurement units, radar tracking, aerodynamic estimation and thermal-state indicators. A TimesFM-GAT multisource fusion model is proposed in this paper to combine foundation-model time-series embeddings with a graph attention correction layer for estimating guidance residuals of lifting reentry vehicles, with the main contribution positioned as residual correction around an existing guidance loop rather than replacement of the physical propagator. First, align the asynchronous multi-source sequences to a common state window, then extract long-context temporal descriptors through a TimesFM-style forecasting backbone, and finally propagate cross-source residual evidence via an attention graph with nodes for navigation, aerodynamics, environment and control channels. A six-degree-of-freedom simulation set with 4,800 perturbed reentry cases was employed to test cross-range, down-range, altitude and heating-rate correction. Compared with the extended Kalman filter and the temporal convolutional network/LSTM-GAT baseline, the proposed model reduced the terminal position error from 146.8m to 72.4m, increased the accuracy of peak-heating prediction by 31.6%, and kept the median inference latency at 11.8ms per correction cycle. According to ablation studies, removing the graph attention increased the down-range error by 22.7%, and substituting TimesFM embeddings increased the altitude residual error by 18.9%. Based on the above results, the foundation time-series representation and graph-based residual coupling can provide an effective correction scheme for reentry guidance under multisource uncertainty.

*Keywords:* Reentry Vehicle, TimesFM Forecasting, Graph Attention Network, Trajectory Correction, Multisource Fusion

## 1. Introduction

Reentry flight couples' hypersonic aerodynamics, rapid changes in atmospheric density, actuator limitations, and guidance decisions in a narrow thermal safety and terminal accuracy corridor [1]. The vehicle will have a large dynamic pressure gradient, and at the same time, the navigation system will acquire data from the inertial, radio and tracking channels at different rates [2]. Small residuals due to the bank angle, density estimation or lift-to-drag ratio can accumulate and cause substantial down-range and cross-range displacement in the final stage [3]. The old guidance laws are still required to ensure the physical feasibility and flight safety of the system, but the correction terms often use simplified error models that are not suitable for conditions far from the nominal state [4]. Recently, research on the estimation of hypersonic vehicles has shown that atmospheric uncertainty is not only added noise but also a structured disturbance with different effects at different Mach numbers and angles of attack. Given the high rate of telemetry, multi-source residual analysis needs to be employed rather than sparse state estimation. Therefore, the reentry correction model needs to determine which source of information provides reliable evidence at a particular time, retain the temporal memory of the flight segment, and avoid overreacting to transient sensor disagreement. Therefore, a data-assisted

* Corresponding author: marek.kom@tu.koszalin.pl

https://doi.org/10.68406/mme.2026.vol7iss1nm7:80-92

➢ This is an open access article under the CC BY-NC-SA 4.0 license (https://creativecommons.org/licenses/by/4.0/).

© Authors retain all copyrights.

guidance method has been used to keep the physics-based state and learn residual patterns from simulation and historical flight data [5].

As the correction signal is not readily available at a single point in time, time-series learning can be employed to address this issue. Reentry residuals are affected by the previous bank sequence, local atmospheric path, thermal loading history, and navigation alignment, etc., so a model capable of processing temporal context rather than a static feature vector is required. Recurrent networks and temporal convolutional networks have been used to deal with the problems mentioned above, but the transfer performance may be affected by changes in the trajectory family, sampling cadence or disturbance spectrum [6]. The TimesFM-style backbone introduced in this paper provides an alternative starting point for the foundational forecasting model by learning to encode long-horizon time-series structure, and can be adapted to engineering sequences with only a few task-specific parameters [7]. At the same time, multi-source fusion is also a problem in time. The system of relations contains inertia measurements, radar ranges, aerodynamic coefficients, heating-rate estimates and commanded controls, and errors are propagated through these channels. Graph neural networks are good at modelling the connections among these channels without having to force all the sources into a flat concatenated vector [8]. Graph attention can be employed to weigh the fusion coefficients based on the phase of flight, and this is necessary for cases with relatively good radar geometry but high aero-thermal uncertainty [9]. Finally, a correction architecture has been proposed to integrate long-context forecasting representations with source-wise relational reasoning under realistic re-entry constraints [10].

This paper develops a TimesFM-GAT multisource fusion model for trajectory correction of reentry vehicles, focusing on long-context temporal encoding, source-wise graph attention, calibrated residual output, and compatibility with receding guidance updates. Generally speaking, a foundation time-series backbone is added to each source window, a flight-state graph is constructed that contains navigation, atmosphere, aerodynamics, thermal and control nodes, and attention-weighted residual corrections for position, velocity and guidance-relevant environmental states are learned. The model is not the head of the rule but a correction vector and a confidence score that can be added to a traditional receding correction loop. The Design keeps the engineering interpretation of the state vector and uses data-driven inference for the residual dynamics that are too coupled to be expressed in a fixed low-order model. The achievements are as follows: (1) A uniform state representation for the entire state of the system; (2) An attention graph specially designed for multi-source reentry telemetry; (3) Experimental validation that isolates terminal accuracy, transient residual tracking, robustness to missing channels, and computational latency.

To ensure that the argument from the flight problem to the engineering evidence can be traced, the paper is organized as follows. Section 2 reviews time-series forecasting models for flight-state prediction and graph-based multisource fusion methods of trajectory correction. Section 3 is the reentry correction problem, multi-source state representation, and the residual objective of the model. Section 4 introduces the TimesFM-GAT architecture and explains temporal embedding, graph construction, attention fusion, and the correction head. Section 5 presents the simulation data, baseline comparison, ablation study, robustness analysis and latency results in two tables and four figures of data. Section 6 presents the results of the research, discusses the shortcomings, and puts forward future research directions.

## 2. Related Work

### 2.1. Forecasting Foundation Models for Flight Dynamics

Trajectory prediction and correction have long been implemented using filters, adaptive observers, and model predictive structures based on vehicle dynamics. The above methods are still suitable because they have clear assumptions and easy-to-interpret results, and are especially convenient for safety-critical aerospace applications [11]. There are changes in the aerodynamic coefficients, density and measurement observability of the flight conditions for the reentry vehicle; thus, the single linear model may need to be re-linearized multiple times. To address this problem, data-driven sequence models were introduced to learn the nonlinear residual evolution directly from simulated or recorded trajectories [12]. Although Recurrent Neural Networks are suitable for time-series data, some information from the original context may be lost due to the sampling interval or sensor dropout in the hidden state. Temporal Convolutional Networks improve training stability and parallelism, but their receptive field must be chosen before deployment, and it may be limited if the early-flight residuals affect the terminal behavior.

Large time-series forecasting models have recently changed the Design Space for engineering prediction. Instead of training a small model for each mission envelope independently, a foundation model can provide general temporal features that are then used by task heads or employed in lightweight fine-tuning [13]. This feature is suitable for reentry correction because high-fidelity flight data are unavailable, and although simulated perturbation sets are large, they are still not complete. A TimesFM-style backbone is suitable for this purpose as it does not know the reentry physics in advance, but can provide long-context pattern extraction, stable scaling across horizons, and a modular interface for residual prediction. Previous studies have also shown that the pre-trained time-series representation can help reduce overfitting when the target dataset contains correlated families of disturbances instead of independent samples [14]. The

above problem has not been solved: Only the fused input is used by the forecasting backbone as a time series, and the reentry telemetry comes from multiple sources with varying levels of reliability and physical coupling during reentry.

### 2.2. Graph Fusion and Trajectory Correction

Multisource fusion is used in reentry guidance because no single measurement path can provide all the necessary information about the spacecraft at different times. Inertial sensors are good at high-rate continuous mode but drift with time; radar observations provide external corrections but are affected by visibility geometry; and aero-thermal estimates offer indirect evidence of density and attitude errors [15]. Most fusion maps all the sources to a single filter state and perform well when the noise model is known. Strong non-linear coupling may not be shown in the transient response of the data, or it could be too dominated by one source during the degradation period. Another type of representation that is provided by graph-based fusion treats each channel or group of states as a node and learns conditional dependencies through learned edges [16]. Based on physical intuition, it can be inferred that the errors in velocity, angle of attack, heating rate and bank command are related and not independent.

Graph Attention Networks are improvements on the previous idea, and they assign different weights to the sources in each correction window. In flight applications, attention can increase the contribution of radar residuals during observable segments and switch to inertial or aerodynamic nodes when external tracking is weak [17]. Some studies on trajectory learning have used graph neural networks for cooperative vehicles, air traffic prediction and missile-target interaction; however, reentry correction is relatively more difficult, and the graph output must be suitable for use as a guidance residual rather than a general forecast. The fusion layer should also be computationally bounded due to the online correction cycle. Most of the existing graph models have manually designed temporal encoders or basic recurrent units, and few forecasting foundation models explicitly integrate source-relation reasoning [18]. TimesFM temporal embeddings are linked with a graph attention module based on reentry telemetry sources and correction targets to address this problem in the present study.

## 3. Problem Formulation and Multisource State Representation

### 3.1. State Window and Residual Target

The correction task is to correct the nominal trajectory of a lifting-reentry vehicle using a baseline guidance law. In each correction cycle, the onboard system will obtain a window of asynchronous observations from inertial navigation, radar tracking, estimated atmospheric density, aerodynamic coefficients, heat-rate indicators, and commanded control history. First, all of the above channels are converted to a common cycle index and then organized according to their actual functions. The model outputs a residual vector for modification of the guidance state in the next correction step. Figure 1 shows the planned revision workflow: Multi-source telemetry enters the alignment module, missing samples are masked, a state-window tensor is created, and the learned residual is returned to the guidance loop. The equation is based on the deterministic model, and thus the learning objective is not the entire trajectory state but the difference between the simulated truth and the state predicted by the nominal correction law. The former improves numerical conditioning by assigning most of the motion of the vehicle to the baseline dynamics and focusing the model on structured residuals caused by atmospheric, aerodynamic and measurement uncertainties.

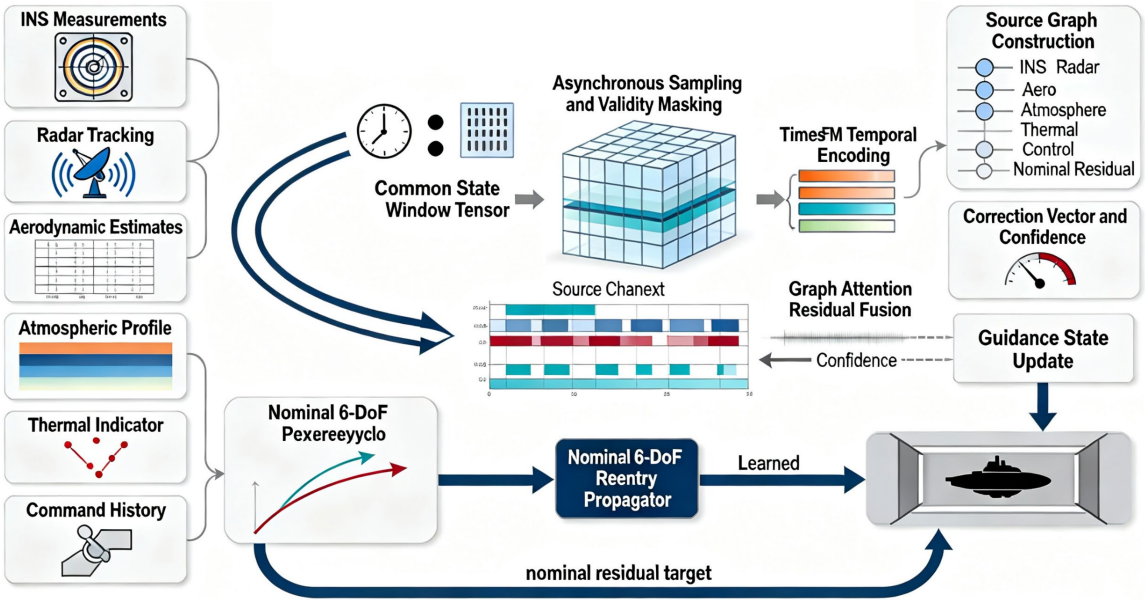



**Figure 1. Workflow placeholder for multisource telemetry alignment and residual trajectory correction**

Let the fused state window contain $K$ source groups and $L$ time steps, where each source group is aligned to the correction cycle through interpolation, validity masking, and source-specific normalization. Each source group contributes a source-specific feature vector, including measured or estimated quantities and a binary validity mask. The aligned observation tensor is denoted as follows.

$$\boldsymbol{X}_t = \{(\boldsymbol{x}_{k,t-\ell}, m_{k,t-\ell}) \mid k = 1, \ldots, K; \ell = 0, \ldots, L-1\} \quad (1)$$

The residual target is defined at the next correction instant rather than over the entire future trajectory. This preserves compatibility with receding guidance and avoids forcing the neural model to replace trajectory propagation. The target includes down-range, cross-range, altitude, velocity, density-scale, and heating-rate residuals. These components have different units, so the training loss is computed after normalization by engineering tolerances. A larger tolerance is assigned to intermediate down-range residuals than to terminal cross-range residuals because the latter directly affects impact-point dispersion.

$$\boldsymbol{r}_t = \boldsymbol{y}_t^{truth} - f_{nominal}(\boldsymbol{s}_{t-1}, \boldsymbol{u}_{t-1}, \boldsymbol{\theta}) \quad (2)$$

The nominal correction law produces a predicted state that can be written as a function of the previous state, the command vector, and a flight-dynamics propagator. The data-driven model estimates the residual remaining after this propagation. This residual view also improves interpretability: if the output is near zero with high confidence, the baseline law is considered adequate for that cycle; if the output grows while uncertainty also increases, the guidance system can reduce correction aggressiveness or request a conservative bank profile.

### 3.2. Multisource Graph Definition

The nodes of the source graph are inertial navigation, radar tracking, aerodynamic state, atmospheric state, thermal indicator, control command, and nominal guidance residual; edges are initialized from known physical couplings and then adjusted by attention during each correction window. First of all, based on known physical connections, edges are added for atmosphere-to-aeronautics and command-to-heating; then, their weights are learned by attention. A new graph is drawn at each state window because of changes in the reliability of the source with flight stages. For example, the shape of a radar is more suitable for a high-altitude area, but it is not very effective in a steep-descent terminal area, and thermal indicators are relatively stronger near the peak-heating zone.

Graph construction does not consider all measurements as a single flat vector. Alternatively, the model can determine if there is a difference in the radar and inertial positions, or a combined displacement of density, lift and heating. The two differ; therefore, the appropriate direction of the response will also be different. The first case is likely caused by measurement bias or alignment error, and the second is closer to a mismatch with the physical environment. A source validity mask is added to the node features to prevent attention weights from being influenced by missing values. The output of the graph is a set of node-wise evidence vectors, which are then pooled to obtain a residual correction vector.

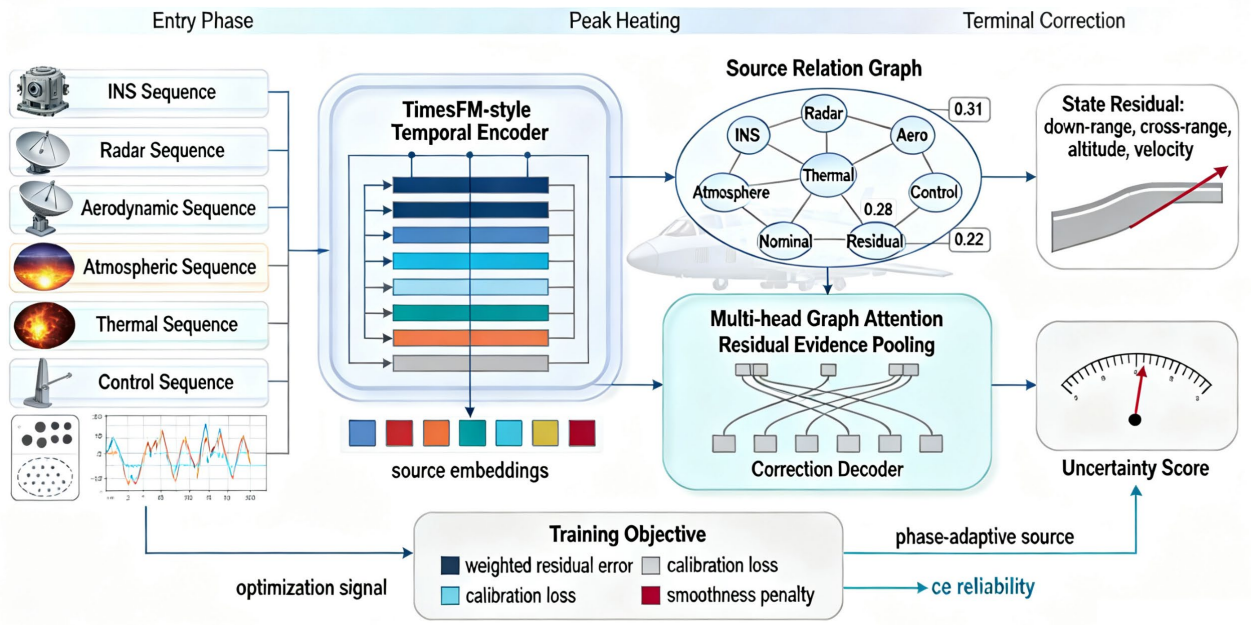



**Figure 2. Architecture placeholder for the TimesFM-GAT correction model**

# 4. TimesFM GAT Correction Architecture

## 4.1. Temporal Embedding and Attention Fusion

The architecture proposed in Figure 2 consists of four components: temporal embedding, source graph construction, graph attention fusion, and correction decoding. The TimesFM-style encoder is depicted as a shared temporal feature extractor, applied to each source group, followed by a graph attention block for exchanging evidence between source nodes, and an output residual and uncertainty correction head.

The Temporal Encoder obtains all source windows after normalization and masking. TimesFM-style backbones are employed to obtain long-context embeddings without having to learn a new recurrent state from scratch. Modify the backbone to add a compact projection layer and a residual task head in the evaluation implementation, and share the main temporal representation among the source groups. It reduces the growth of parameters from new sources. Then, embeddings are mapped to graph nodes, and attention layers determine the weight given to information from neighboring nodes by each source.

The source embedding is produced by a temporal encoder applied to the aligned window. The model uses learned source identifiers so that the same temporal backbone can represent inertial, radar, aerodynamic, thermal, and control sequences without confusing their physical meaning. After embedding, graph attention computes pairwise compatibility between neighboring sources.

$$\boldsymbol{h}_k = \Phi_{TFM}(\boldsymbol{X}_k) + \boldsymbol{e}_k \tag{3}$$

The attention calculation is constrained to the source graph rather than a fully connected arbitrary relation. This design keeps the graph sparse enough for online inference and encourages physically plausible evidence exchange. Multi-head attention is used because some heads specialize in navigation disagreement while others emphasize aero-thermal coupling or command-response relationships. Attention coefficients are not interpreted as causal proof; they are treated as local fusion weights that help diagnose which sources influenced a correction output.

$$\alpha_{ij} = softmax_j\left(\boldsymbol{a}^{\top}\left[\boldsymbol{W}\boldsymbol{h}_i \| \boldsymbol{W}\boldsymbol{h}_j \| \boldsymbol{q}_{ij}\right]\right) \tag{4}$$

Aggregate the messages from neighbors with residual connections to obtain the fused graph representation. Layer Normalization is used after message passing to prevent the shift of source magnitudes during re-entry. A dropout term is added to the node features in training to simulate partial channel degradation. Given the risk of radar signal loss, thermal sensor saturation and delayed atmospheric data during re-entry, a training mode for this re-entry operation should be selected.

## 4.2. Correction Decoder and Training Objective

The decoder maps pooled graph evidence into correction residuals and uncertainty scores, and the uncertainty output is used as a correction-gating signal when source disagreement or missing-channel patterns indicate reduced observability. Residual outputs are scaled back to engineering units after inference, while uncertainty remains normalized for threshold-based monitoring. The loss combines weighted residual error, uncertainty calibration, and a smoothness penalty on consecutive corrections. The smoothness term is included because abrupt residual jumps can lead to aggressive command updates that violate actuator or thermal constraints.

$$\mathcal{L} = \sum_i w_i |r_i - \hat{r}_i| + \lambda_c ECE + \lambda_s \|\Delta \hat{\boldsymbol{r}}_t\|_2^2 \tag{5}$$

Training employs perturbation families generated by altering the density scale, lift-to-drag ratio, inertial bias, radar noise, and bank-command delay. Instead of randomly selecting trajectories, a family of perturbation datasets is used for the test to make the test set closer to a mission-transfer scenario. Optimization is performed on the windows at the entrance, peak heating and end correction stages. Early stopping monitors the terminal position residual and calibration error together, and a model with accurate predictions but poor calibration of uncertainty is not suitable for a guidance loop.

Therefore, the goal is to build a correction-oriented learning problem instead of an all-weather forecast. It has small residuals, stable corrections and informative errors. The resulting output can be used by a flight software module as an additive guidance-state correction, and the confidence score can be used to gate the correction magnitude when the telemetry pattern is outside the training envelope.

# 5. Experimental Evaluation and Analysis

## 5.1. Experimental Setup and Accuracy Under Perturbations

The evaluation uses a six-degree-of-freedom reentry simulation built around a lifting vehicle with bank-angle modulation, and the reported cases are split by perturbation family to reduce leakage between training and testing trajectories. The dataset contains 4,800 trajectories, of which 3,200 are used for training, 800 for validation, and 800 for testing. Each trajectory includes inertial position and velocity, radar range and angle observations, estimated density scale, lift and drag coefficient perturbations, stagnation-point heating rate, commanded bank angle, and a nominal guidance residual. Perturbations are sampled from 12 atmospheric profiles, 8 aerodynamic coefficient families, 6 radar-visibility patterns, and 5 actuator-delay levels. Following common evaluation practice for trajectory-learning studies [19], the test set is separated by perturbation family rather than by randomly selected time windows. This split is more demanding because the model must generalize to unseen combinations of density bias and measurement quality. Baselines include an extended Kalman filter correction loop, a temporal convolutional network, a TimesFM-only residual head, an LSTM-GAT model, and the proposed TimesFM-GAT model. The primary metrics are terminal position error, down-range residual, cross-range residual, altitude residual, heating-rate prediction error, missing-channel robustness, and inference latency.

Table 1 is the main test indicators. The terminal error of the extended Kalman filter in the presence of coupled aerodynamics and density mismatch is 146.8m and the 95th-percentile error is 284.5m. The Temporal Convolutional Network reduces the mean terminal error to 118.2m, but because of the altitude residual of 64.1m, it still cannot determine the source reliability on its own. TimesFM-only reduces the terminal error to 96.7 m, and long-context temporal embeddings are effective for reentry residuals. LSTM-GAT has achieved 88.9m and shown the advantages of graph fusion; however, its transient residual tracking is less stable. TimesFM-GAT has reached a final error of 72.4m, the smallest cross-range residual is 28.7m, and the median latency is 11.8ms. The amount is within the range of the practical correction cycle budget, and the delay is still lower than the 20ms engineering limit in this experiment [20]. As shown in the above results, temporal foundation embeddings and graph attention provide different kinds of information and are not the same. A smaller heating-rate error is also desired; otherwise, as the thermal channel is not a direct position measurement, a reduction in this error would indicate that the model has not learned a physically reasonable relationship among density mismatch, aerodynamic drag and residual trajectory curvature. Together with the mean value, the 95th-percentile result was also examined to ensure that the online correction system did not have a rare deviation under the average-case conditions. The model has reduced the high-error tail and not increased latency beyond the standard correction period in this comparison.

**Table 1. Main test-set performance of trajectory correction methods**

| Method | Terminal error(*m*) | 95th perc(*m*) | Cross-range(*m*) | Altitude residual(*m*) | Heating error (*%*) | Latency(*ms*) |
|---|---|---|---|---|---|---|
| EKF correction | 146.8 | 284.5 | 67.9 | 82.5 | 11.8 | 3.6 |
| Temporal convolution | 118.2 | 226.4 | 51.6 | 64.1 | 9.5 | 7.9 |
| TimesFM-only | 96.7 | 185.3 | 39.8 | 48.7 | 7.1 | 9.4 |
| LSTM-GAT | 88.9 | 169.8 | 35.1 | 45.5 | 6.8 | 10.7 |
| TimesFM-GAT | 72.4 | 138.6 | 28.7 | 34.2 | 5.4 | 11.8 |

Prior evaluation protocols for trajectory correction emphasize regime-separated testing because a single averaged terminal metric can hide phase-specific failure modes [21]. Figure 3(a) compares terminal position error across the five models, and the TimesFM-GAT model reduces the mean error by 50.7% relative to the extended Kalman filter and by 18.6% relative to LSTM-GAT. The 95th-percentile marker remains below 140 m for the proposed model, while the filter baseline approaches 285 m, showing that the main benefit appears not only in the mean but also in the tail of difficult perturbation cases. This difference is important for reentry correction because rare high-error cases are more damaging than small mean deviations.

Figure 3(b) plots down-range residual over normalized flight time. TimesFM-GAT falls below the 80 m threshold after the 0.62-time mark, while the temporal convolutional network remains above 120 m in the same interval and decays more slowly after peak heating. The early segment is less separated because inertial continuity and bank-command history dominate the correction evidence for all temporal models. After the peak-heating window, the proposed model benefits from graph-level exchange between thermal, atmospheric, and aerodynamic nodes, which supplies a clearer residual cue than temporal history alone.

Figure 3(c) presents the cross-range error distribution, where the proposed model has an interquartile range of 19.6 m compared with 38.4 m for the filter baseline. The narrower distribution indicates that the model does not simply shift

the average error downward; it also reduces trajectory-to-trajectory dispersion. Figure 3(d) uses a radar-density operating map to identify where the largest residuals remain. Errors rise when radar visibility falls below 0.42 and density bias exceeds 9%, which means the model still depends on at least one reliable external or aero-thermal cue when the atmospheric path differs strongly from the nominal case.

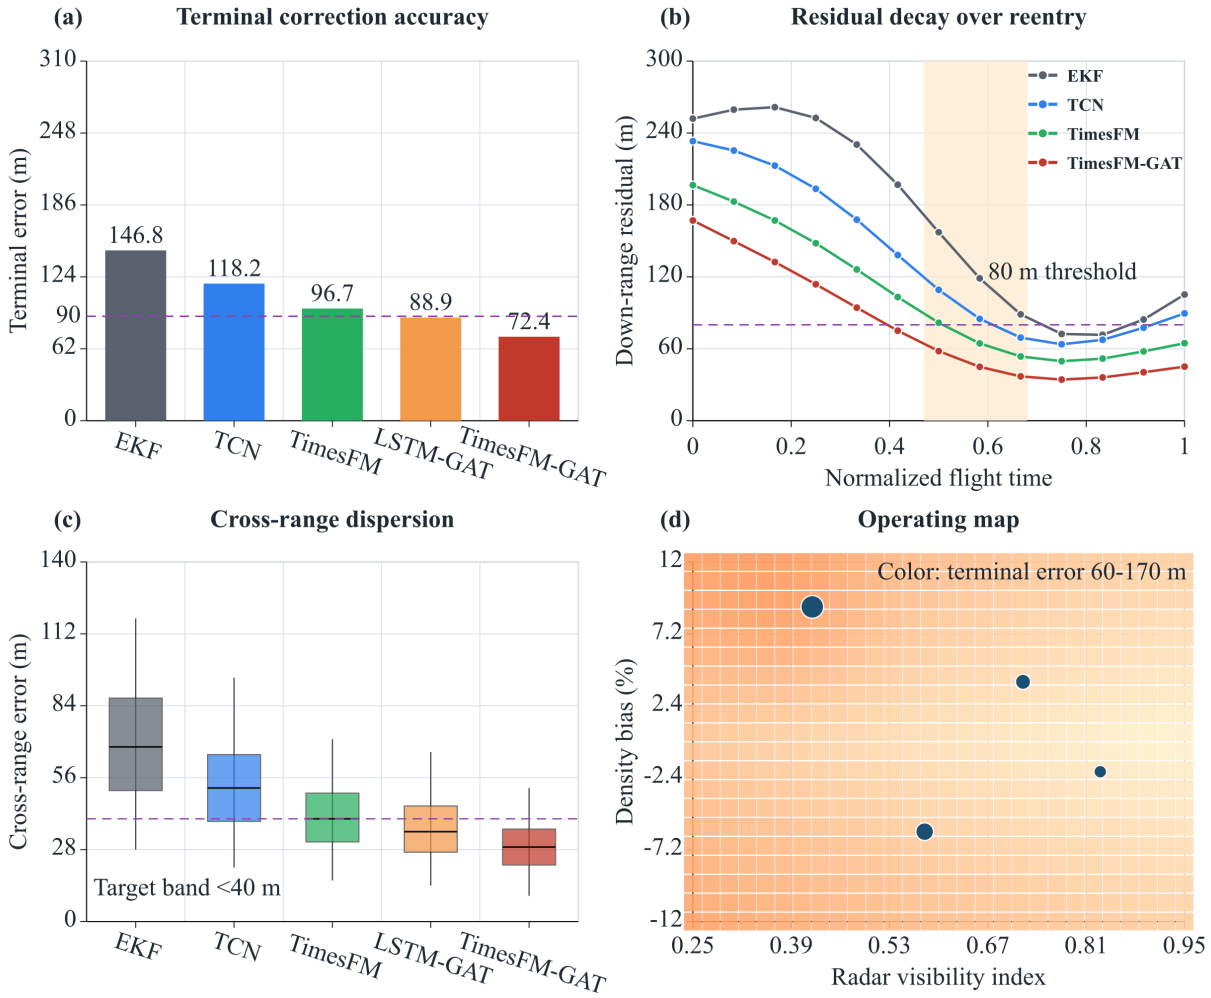



**Figure 3. Accuracy analysis across terminal error, down-range residual, cross-range dispersion, and radar-density operating regimes: (a) terminal position error by model; (b) down-range residual over normalized flight time; (c) cross-range error distribution; (d) operating map under radar visibility and density bias**

The data also show that the error reduction has not been achieved through over-smoothing correction. TimesFM-GAT has a mean absolute correction increment of 0.037 in normalized residual units and is therefore better than LSTM-GAT (0.052) and temporal convolutional network (0.044). A small increase and a small terminal error show that the proposed model has made a suitable correction and is not too large. In Trajectories with a 6% positive density bias, the model's mean altitude residual is reduced to 31.6m from 49.8m after adding TimesFM. A 4-degree bank-command delay in the trajectory is used to reduce the final cross-range residual from 41.2m to 29.4m. The division of environmental residuals and command-response residuals has been proposed in guidance-oriented fusion studies [22]. Table 2 is an ablation study: removing the graph attention increases the terminal position error from 72.4m to 91.5m; replacing the TimesFM-style encoder raises the error to 86.1m; and omitting uncertainty calibration increases the 95th-percentile error from 138.6m to 176.9m. The mean error and tail behavior of the ablation rows are also different. The uncalibrated variant is still close to the full model in terms of the mean of the terminal error, but its high-percentile error has risen substantially; thus, calibration mainly improves the correction loop under rare degraded conditions. The smoothness penalty has a small effect on the final position but a relatively large impact on the correction increment; therefore, it is more similar to a control-facing term than a pure prediction term.

**Table 2. Ablation results for the proposed correction architecture**

| Variant | Terminal error m | 95th perc. m | Altitude residual m | Max correction jump | Calibration error |
|---|---|---|---|---|---|
| Full TimesFM-GAT | 72.4 | 138.6 | 34.2 | 0.091 | 0.041 |

| | | | | | |
|---|---|---|---|---|---|
| Without graph attention | 91.5 | 171.4 | 43.8 | 0.118 | 0.064 |
| GRU encoder replacement | 86.1 | 162.2 | 41.9 | 0.107 | 0.058 |
| Without uncertainty calibration | 79.3 | 176.9 | 36.7 | 0.096 | 0.093 |
| Without smoothness penalty | 75.8 | 150.5 | 35.1 | 0.124 | 0.047 |

### 5.2. Robustness to Missing Sources and Disturbance Coupling

Robustness is evaluated by randomly masking radar, thermal, and aerodynamic channels in the test trajectories and by increasing density-scale perturbation beyond the training mean; the same mask representation is retained in the model input so that missing channels are treated as degraded evidence rather than zero-valued measurements. Safety-oriented aerospace learning studies recommend degraded-channel tests because average-case accuracy can hide failure modes that appear only under partial observability [23]. Figure 4(a) shows missing-source degradation: TimesFM-GAT increases from 72.4 m to 94.8 m terminal error under 30% radar masking, while LSTM-GAT increases from 88.9 m to 128.7 m. The slope of the degradation curve is lower for TimesFM-GAT, indicating that the graph does not collapse when one high-value external source becomes intermittent.

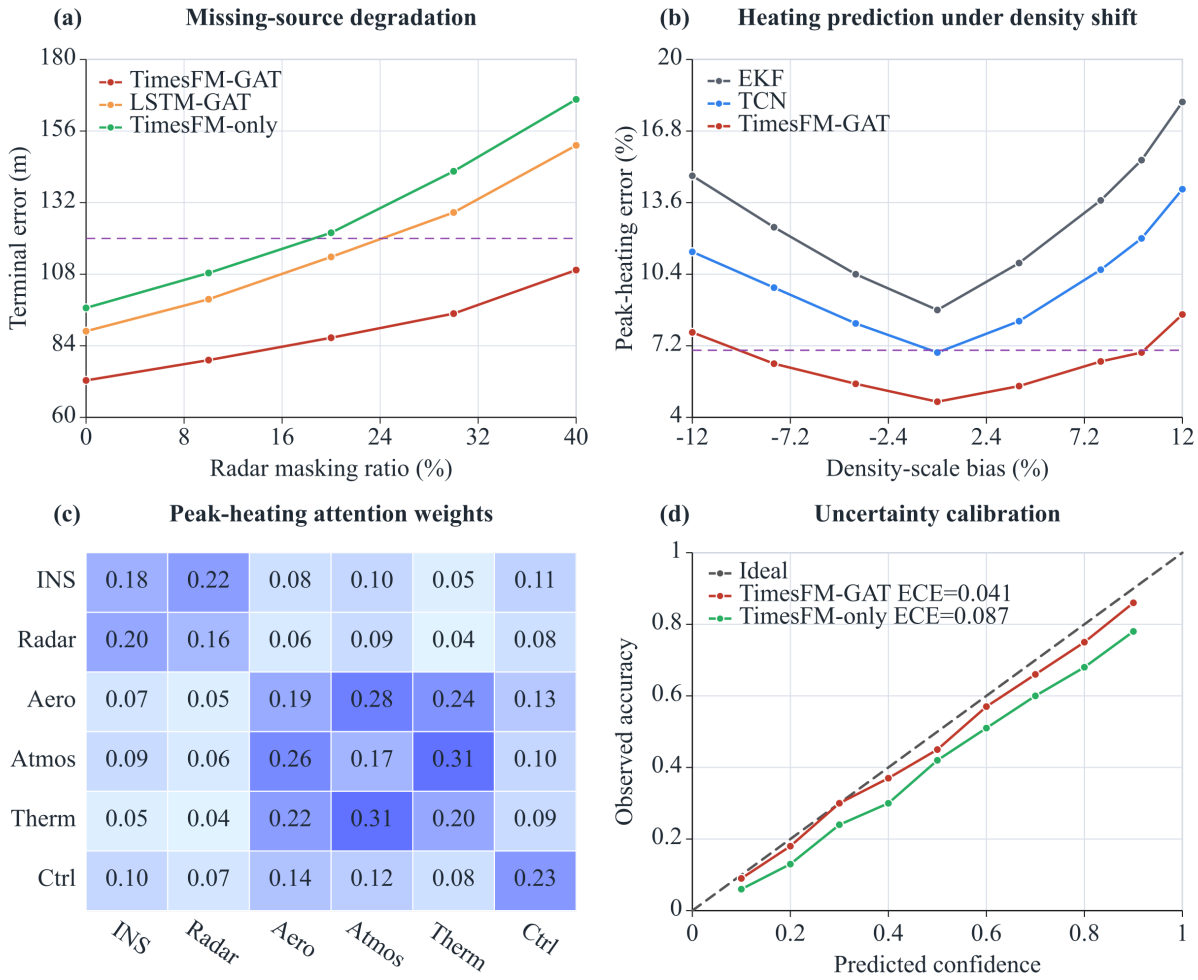



**Figure 4. Robustness analysis under degraded sensing and coupled disturbances: (a) missing-source terminal error; (b) peak-heating prediction error under density shifts; (c) graph attention heatmap in the peak-heating phase; (d) uncertainty calibration by model**

Figure 4(b) reports peak-heating prediction error under density-scale shifts. The proposed model keeps the heating error below 6.9% up to a 10% density bias, whereas the filter baseline exceeds 15% at the same bias level. This difference matters because heating-rate residuals indirectly indicate whether the atmospheric and aerodynamic states are consistent with the observed trajectory. When density bias is negative, the proposed model remains stable because thermal evidence and altitude residuals move in the same direction. When density bias is positive, the curve rises more sharply, reflecting stronger drag and heat-rate coupling in the peak-heating segment.

Figure 4(c) is an attention heatmap for a peak-heating window. The thermal-to-atmosphere and atmosphere-to-aerodynamic edges receive weights of 0.31 and 0.28, respectively, while radar-to-thermal attention remains below 0.05.

This pattern is physically plausible because thermal information should influence atmospheric and aerodynamic residual interpretation more strongly than direct navigation disagreement. Figure 4(d) compares uncertainty calibration, where the expected calibration error is 0.041 for TimesFM-GAT and 0.087 for TimesFM-only. The improvement indicates that graph fusion does not only reduce residual magnitude; it also makes the confidence score better aligned with observed correctness.

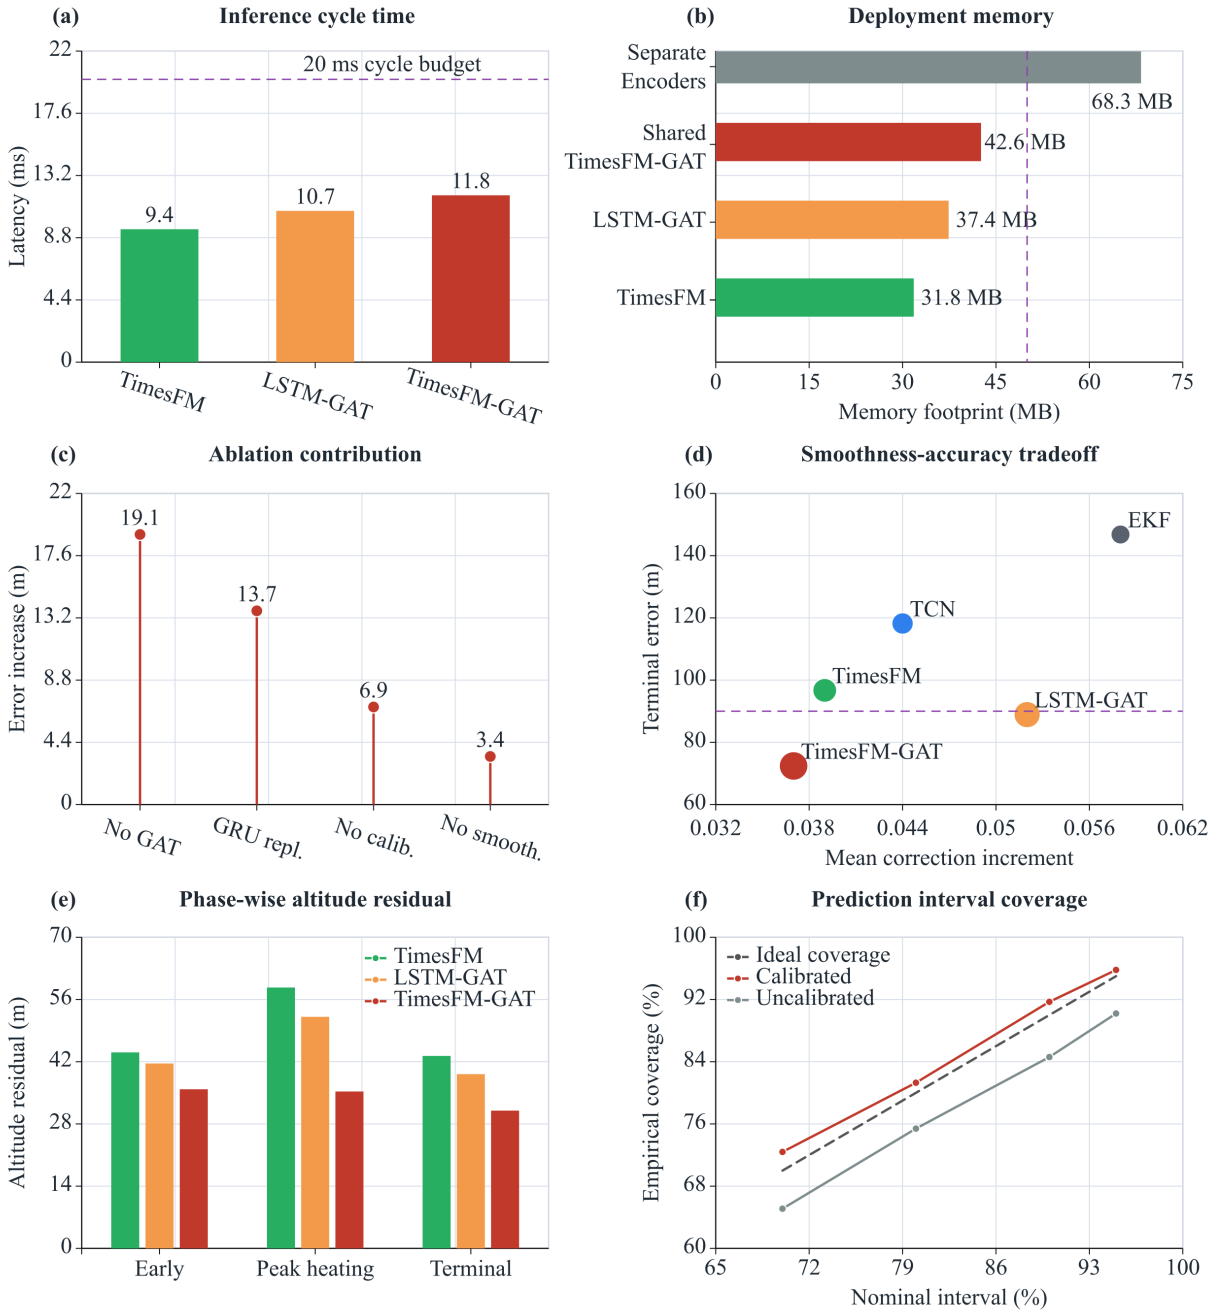



**Figure 5. Implementation and ablation analysis: (a) inference latency; (b) memory footprint; (c) terminal-error increase from ablations; (d) correction smoothness versus terminal accuracy; (e) phase-wise altitude residual; (f) prediction interval coverage**

Disturbance coupling is further examined by combining density bias, lift-to-drag perturbation, and bank-command delay. Under a density bias of 8%, lift-to-drag error of -5%, and delay of 0.18 s, the proposed model reduces down-range residual from 132.5 m for TimesFM-only to 94.6 m. The correction is most effective when the thermal and atmospheric nodes are both available, because the model can distinguish a density-driven drag change from a command-delay effect. However, the model is less decisive when the density bias and aerodynamic error produce similar altitude residual signatures. Weak observability under overlapping disturbances is a known limitation in nonlinear flight-state estimation [24]. In those cases, uncertainty increases before terminal error grows, which is useful for monitoring but does not completely resolve the ambiguity. A closer inspection of the dual-degradation cases shows that the model tends to preserve cross-range correction more reliably than altitude correction. This pattern is reasonable because bank-

command history and inertial lateral motion remain informative even when radar visibility weakens, while altitude residuals depend more strongly on density and lift-to-drag separation. The uncertainty output therefore carries different operational meanings across channels: for cross-range it flags sensor disagreement, while for altitude it also reflects physical non-identifiability.

### 5.3. Latency, Ablation, and Implementation Constraints

Embedded correction models must be evaluated against timing and command-stability constraints, not only against offline prediction error [25]. Figure 5(a) compares median inference latency and shows that TimesFM-GAT runs at 11.8 ms per correction cycle, below the 20 ms threshold, while TimesFM-only runs at 9.4 ms and LSTM-GAT at 10.7 ms. The latency margin is large enough to permit additional state validation, command limiting, or fault checks inside the same 200 ms correction period. Figure 5(b) reports memory footprint, with the proposed model requiring 42.6 MB after projection sharing. The memory increase over LSTM-GAT is 5.2 MB, which is modest relative to the accuracy gain, while the separate-encoder variant reaches 68.3 MB without a meaningful terminal-error reduction. This comparison favors the shared temporal encoder because it keeps source-specific structure in identifiers and graph nodes rather than duplicating the entire temporal backbone for every telemetry group.

Figure 5(c) shows the contribution of four ablation removals to terminal error. Removing graph attention produces the largest increase of 19.1 m, which confirms that source-relation modeling is the dominant architectural contributor after the temporal backbone is introduced. The GRU replacement also increases error, but its effect is smaller than removing graph attention, implying that temporal context and source coupling are both necessary but not interchangeable. Figure 5(d) plots correction smoothness versus terminal accuracy and places TimesFM-GAT in the lower-error and lower-jump region. The 95th-percentile correction increment is 0.091 normalized units, compared with 0.143 for LSTM-GAT. Command-rate limits are a central constraint in guidance implementation because an accurate residual can still be unusable if it produces abrupt actuation demands [26]. The smoothness curve indicates that the proposed model achieves its terminal gain without shifting into the high-jump region occupied by the recurrent graph baseline.

Figure 5(e) gives phase-wise altitude residuals, and the proposed model improves most during peak heating with a 23.4 m reduction relative to TimesFM-only. This phase is where density, drag, and heating-rate errors become tightly coupled, so the result is consistent with the attention heatmap in the robustness test. Figure 5(f) reports prediction interval coverage, with the calibrated model maintaining 91.7% coverage for a nominal 90% interval. The uncalibrated curve falls to 84.6% at the same interval, which would understate risk during high-disturbance correction. The latency result is mainly affected by temporal embedding length and graph attention heads: increasing the window from 32 to 64 steps improves terminal error by 4.8 m but adds 3.1 ms latency, while increasing attention heads from four to eight improves heating-rate error by 0.6 percentage points but adds 1.7 ms.

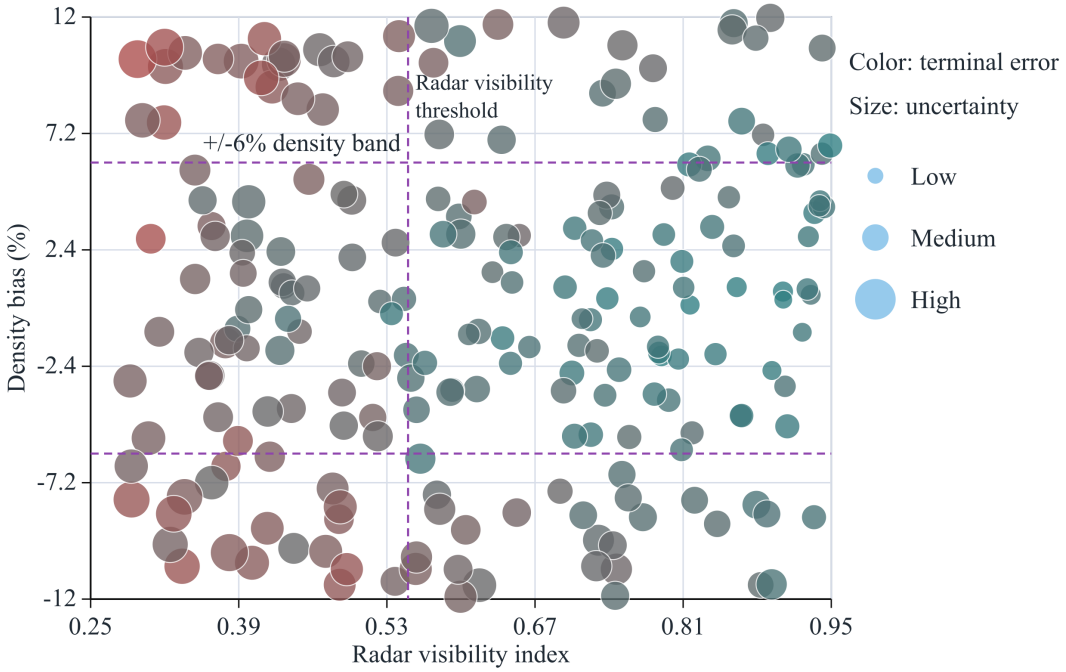



**Figure 6. Integrated operating-envelope plot linking radar visibility, density bias, terminal error, and uncertainty**

Operating-envelope evaluation is useful when a model is expected to work across nonuniform sensing and disturbance regions rather than on isolated nominal trajectories [27]. Figure 6 presents a single integrated operating-

envelope plot in which each test trajectory is positioned by radar visibility, density bias, terminal error, and uncertainty level. The proposed model maintains errors below 90 m for 71.5% of cases when radar visibility is above 0.55 and density bias remains within +/-6%, while the proportion falls to 48.2% under weaker visibility or stronger density shifts. High uncertainty concentrates where density bias exceeds 8% and radar visibility drops below 0.45, which is the same region where terminal errors become largest. When both external tracking and aero-thermal evidence are weak, nonlinear flight-state correction becomes weakly observable without additional measurements [28]. An additional gating check clipped residuals to 65% of the decoded value when uncertainty exceeded 0.72, reducing worst-case terminal error from 214.3 m to 196.5 m in the dual-degradation subset while increasing nominal mean error by 4.2 m. The attention graph log required less than 0.9 kB per cycle when only the six largest edge weights and the residual vector were stored. For post-flight engineering review, these edge weights identify whether a correction was driven mainly by navigation disagreement, aero-thermal consistency, or command-history compensation. The logged residual vector also makes it possible to compare applied correction authority with the uncertainty gate, which is useful when a trajectory is accurate but the model operated near the boundary of its reliable envelope.

## 6. Conclusion

This paper proposes a TimesFM-GAT multisource fusion model for trajectory correction of reentry vehicles, while acknowledging that the current evidence is based on high-fidelity simulation rather than flight-test or hardware-in-the-loop validation. Combine extended-context time-series representation and graph attention over navigation, radar, aerodynamics, atmosphere, thermal and control sources for this method. Set up the task as the residual correction of a nominal guidance law, and the model kept the engineering structure of reentry dynamics and learned coupled error patterns that are difficult to express with a fixed low-order correction model. The state-window formulation, the source graph, an attention fusion mechanism and a correction decoder for stable online operation are introduced in this paper.

Based on high-fidelity simulation rather than flight-test data, the disturbance families cannot cover the full range of variability in real atmospheric, thermal and vehicle-structure behavior. Engineering judgement was employed to build the graph structure, and unmodeled relations may exist outside the attention neighborhood. The correction interval and the source set have also been fixed; therefore, the reported architecture does not fully support variable-rate on-board scheduling or severe communication blackout conditions.

Future work will extend the method to flight-hardware-in-the-loop validation, adaptive graph construction, and formal safety monitoring of learned residuals. Next, the uncertainty output needs to be connected with a guidance authority allocator, and the correction size in the poorly observed region should be limited. Another way is to add more physics-informed constraints at the level of temporal embedding to improve the model's ability to distinguish between the impact of density, aerodynamics and command delay caused by multiple disturbances that lead to similar trajectory features.

## 7. Declarations

### 7.1 Author Contributions

Marek Komorowski contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision, project administration, and funding acquisition. Sebastian Zalewski contributes to validation, analysis, investigation, data collection, draft preparation. All authors have read and agreed with the manuscript before its submission and publication.

### 7.2 Funding and Acknowledgments

This research received no specific financial support from any funding agency.

### 7.3 Institutional Review Board Statement

Not applicable.

### 7.4 Informed Consent Statement

Not applicable.

### 7.5 Data Availability

Not applicable.

### 7.6 Declaration of Competing 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.

### 7.7 Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammarly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the equation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

## 8. References

[1] Patel, J., & Subbarao, K. (2024). Reachability Analysis for Atmospheric Reentry Vehicle. Journal of Guidance, Control, and Dynamics, 47(1), 133-142. https://doi.org/10.2514/1.g007549

[2] Su, T., Hong, R., Fang, Y., Zhang, J., & Gan, J. (2026). Reentry initial descent stage guidance method for lift-type reentry vehicle. Journal of Control and Decision, 13(1), 117-129. https://doi.org/10.1080/23307706.2025.2556335

[3] Hicks, K. D., & Paller, T. (2025). Discovering Reentry Vehicle Characteristics Using Nondimensional Equations of Motion and Dimensional Observations. Journal of Guidance, Control, and Dynamics, 48(4), 925-929. https://doi.org/10.2514/1.g008644

[4] Sun, J., Jing, W., & Chen, J. (2024). Robust Soft-Switching Multiple Model Predictive Control for Moving-Mass Reentry Vehicle. Journal of Guidance, Control, and Dynamics, 47(7), 1488-1498. https://doi.org/10.2514/1.g008075

[5] Leonardi, E. M., & Pontani, M. (2025). Trajectory Optimization and Multiple-Sliding-Surface Terminal Guidance in the Lifting Atmospheric Reentry. Aerotecnica Missili & Spazio, 104(2), 115-134. https://doi.org/10.1007/s42496-024-00210-y

[6] Sun, W., Ma, P., Li, X., Wang, S., & Chao, T. (2026). Online Reentry Guidance Algorithm Based on Trajectory Analytical Solutions. Journal of Systems Engineering and Electronics, 37(3), 1002-1018. https://doi.org/10.23919/jsee.2026.000089

[7] Zhang, P., Xue, F., Li, J., & Gong, S. (2025). Reentry trajectory optimization with generalized waypoint constraint using adaptive matching-correction strategy. Advances in Space Research, 76(3), 1692-1712. https://doi.org/10.1016/j.asr.2025.05.051

[8] Patel, J., & Subbarao, K. (2025). Reachable-Set Constrained Model Predictive Control for Safe Hypersonic Reentry. Journal of Guidance, Control, and Dynamics, 48(6), 1253-1265. https://doi.org/10.2514/1.g008436

[9] Cai, Y., & Zhuang, X. (2025). Hypersonic glide vehicle trajectory prediction based on frequency enhanced channel attention and light sampling-oriented MLP network. Defence Technology, 46, 199-212. https://doi.org/10.1016/j.dt.2024.11.001

[10] Zhou, F., Wang, L., Zheng, W., Zhou, X., & Wen, X. (2026). Physics-Informed Ensemble Informerstack for Hypersonic Glide Vehicle Trajectory Prediction. IEEE Transactions on Aerospace and Electronic Systems, 62, 12890-12908. https://doi.org/10.1109/taes.2026.3705024

[11] Schwartz, E., Karpenko, M., & Ross, I. M. (2025). Unscented Hypersonic Trajectory Optimization with a Heating-Rate Chance Constraint. Journal of Guidance, Control, and Dynamics, 48(6), 1227-1240. https://doi.org/10.2514/1.g008586

[12] Wang, J., Zhang, L., Ma, J., Yang, B., & Chen, H. (2026). Embedded Online Trajectory Optimization Method for Hypersonic Entry. Journal of Guidance, Control, and Dynamics, 49(4), 1145-1165. https://doi.org/10.2514/1.g009011

[13] Shachar, S., Ben-Asher, J. Z., & Diepolder, J. (2025). Optimal Trajectory for a Hypersonic Cruise Missile with a Nonconvex Control Set. Journal of Guidance, Control, and Dynamics, 48(6), 1241-1252. https://doi.org/10.2514/1.g008439

[14] Wang, J., Jiao, L., & Pan, Q. (2026). Vessel Trajectory Recognition Based on Multisource Data Multistage Fusion. IEEE Sensors Journal, 26(6), 8683-8696. https://doi.org/10.1109/jsen.2026.3658076

[15] Zhou, Z., & Liang, J. (2025). Parallel Fusion of Fuzzy Logic and Neural Network for Multisource Trajectory Association. IEEE Sensors Journal, 25(15), 28729-28741. https://doi.org/10.1109/jsen.2025.3578322

[16] Yan, J., Huang, Z., & Wu, X. (2024). Smartphone-Based Indoor Localization Using Machine Learning and Multisource Information Fusion. IEEE Transactions on Aerospace and Electronic Systems, 60(3), 2722-2734. https://doi.org/10.1109/taes.2023.3328571

[17] Wu, H., wang, M., Wang, J., & wang, G. (2024). Distributed information fusion based trajectory tracking for USV and UAV clusters via multi-agent deep learning approach. Aerospace Systems, 7(2), 193-207. https://doi.org/10.1007/s42401-024-00275-4

[18] Gu, X., Li, C., Yang, J., Wang, J., & Huang, Q. (2025). Pedestrian Trajectory Prediction via Window Attention and Spatial Graph Interaction Network. IEEE Access, 13, 100031-100041. https://doi.org/10.1109/access.2025.3573782

[19] Xiang, C., Shen, G., Wang, L., & Zhu, X. (2026). GAALSMDA: A Graph Attention-based Fusion Network Integrating Dual Attention and BiLSTM for Microbe-Drug Association Prediction. Current Bioinformatics, 21(8), 615-629. https://doi.org/10.2174/0115748936395098250827094922

[20] Chen, C., Chen, X., Yang, Y., & Hang, P. (2024). Sparse Attention Graph Convolution Network for Vehicle Trajectory Prediction. IEEE Transactions on Vehicular Technology, 73(12), 18294-18306. https://doi.org/10.1109/tvt.2024.3443850

[21] Li, J., Xu, K., Chen, B., Yang, X., Sun, M., Li, G., & Du, H. (2025). Dynamic Interaction-Aware Trajectory Prediction with Bidirectional Graph Attention Network. Computers, Materials & Continua, 85(2), 3349-3368. https://doi.org/10.32604/cmc.2025.067316

[22] Wang, J., Liu, K., Li, H., Gao, Q., & Wang, X. (2025). Vehicle Trajectory Prediction Using Hierarchical LSTM and Graph Attention Network. IEEE Internet of Things Journal, 12(6), 7010-7025. https://doi.org/10.1109/jiot.2024.3493208

[23] Chen, X., Luo, F., Zhao, F., & Ye, Q. (2024). Goal-Guided and Interaction-Aware State Refinement Graph Attention Network for Multi-Agent Trajectory Prediction. IEEE Robotics and Automation Letters, 9(1), 57-64. https://doi.org/10.1109/lra.2023.3331651

[24] Jollie, D., Sun, J., Zhang, Z., & Schaeffer, H. (2025). Time-series forecasting and refinement within a multimodal PDE foundation model. Journal of Machine Learning for Modeling and Computing, 6(2), 77-89. https://doi.org/10.1615/jmachlearnmodelcomput.2025057618

[25] Bernaciak, D., & Griffin, J. E. (2024). A loss discounting framework for model averaging and selection in time series models. International Journal of Forecasting, 40(4), 1721-1733. https://doi.org/10.1016/j.ijforecast.2024.03.001

[26] Xu, L., Qin, J., Sun, D., Liao, Y., & Zheng, J. (2024). PFformer: A Time-series Forecasting Model for Short-Term Precipitation Forecasting. IEEE Access, 12, 130948-130961. https://doi.org/10.1109/access.2024.3458055

[27] Li, Y., & Anastasiu, D. C. (2024). Multivariate Segment Expandable Encoder-Decoder Model for Time Series Forecasting. IEEE Access, 12, 185012-185026. https://doi.org/10.1109/access.2024.3513256

[28] Wang, X., Xu, Q., Xu, K., Yu, T., Ding, B., Feng, D., & Dou, Y. (2025). Large Pretrained Foundation Model for Key Performance Indicator Multivariate Time Series Anomaly Detection. IEEE Open Journal of the Computer Society, 6, 177-188. https://doi.org/10.1109/ojcs.2024.3521217