

Effective and Efficient Representation Learning for Flight Trajectories

Shuo Liu^{1,2}, Wenbin Li^{2,3}, Di Yao^{2*}, Jingping Bi^{2*}

¹School of Advanced Interdisciplinary Sciences, University of Chinese Academy of Sciences, China

²Institute of Computing Technology, Chinese Academy of Sciences, China

³University of Chinese Academy of Sciences, China

{liushuo22s,liwenbin20z,yaodi,bjp}@ict.ac.cn

Abstract

Flight trajectory data plays a vital role in the traffic management community, especially for downstream tasks such as trajectory prediction, flight recognition, and anomaly detection. Existing works often utilize handcrafted features and design models for different tasks individually, which heavily rely on domain expertise and are hard to extend. We argue that different flight analysis tasks share the same useful features of the trajectory. Jointly learning a unified representation for flight trajectories could be beneficial for improving the performance of various tasks. However, flight trajectory representation learning (TRL) faces two primary challenges, *i.e.* unbalanced behavior density and 3D spatial continuity, which disable recent general TRL methods. In this paper, we propose FLIGHT2VEC, a flight-specific representation learning method to address these challenges. Specifically, a behavior-adaptive patching mechanism is used to inspire the learned representation to pay more attention to behavior-dense segments. Moreover, we introduce a motion trend learning technique that guides the model to memorize not only the precise locations, but also the motion trend to generate better representations. Extensive experimental results demonstrate that FLIGHT2VEC significantly improves performance in downstream tasks such as flight trajectory prediction, flight recognition, and anomaly detection.

Code — <https://github.com/liushuoer/FLIGHT2VEC>

Introduction

With the continual development of air transportation, flights tend to broadcast their current locations for safe and efficient air traffic control (Zhang et al. 2023). The collected flight trajectories are one of the most critical data sources for air traffic management, attracting increasing attention from both academic and industrial fields. Recent works show that deep representation learning has become the dominant technique and achieved significant performance for various flight trajectory-related tasks, such as trajectory prediction (Liu and Hansen 2018; Wu et al. 2022; Zhang et al. 2023; Guo et al. 2024), flight monitoring (Zhang et al. 2016; Fernández et al. 2019), and anomaly detection (Olive and Basora 2020; Memarzadeh, Matthews, and Templin 2022; Memarzadeh, Matthews, and Weckler 2023). However, these works either

*Corresponding authors.

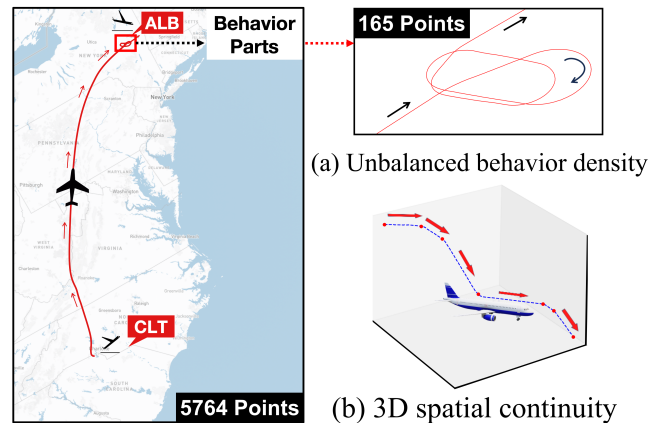


Figure 1: The motivation of FLIGHT2VEC

utilize handcrafted features or design representation models for one specified task. We argue that different flight tasks may share the same useful features of trajectory. Jointly learning a unified representation for flight trajectories could be beneficial for boosting the performance of various tasks, which motivates this work.

Learning unified trajectory representation for different tasks has been well-studied on vehicle and human mobility trajectories (Yao et al. 2017; Li et al. 2018; Chen et al. 2021; Jiang et al. 2023). Deep sequential models, such as Recurrent Neural Networks (Yao et al. 2017; Li et al. 2018) and Transformer (Chen et al. 2021; Yao et al. 2022; Jiang et al. 2023) are employed to encode the spatial-temporal correlations and transform raw trajectories into generic representation vectors. Nevertheless, existing solutions are hard to extend to represent flight trajectories due to the two challenges, *i.e.* unbalanced behavior density and 3D spatial continuity.

Unbalanced behavior density. Flight trajectories are characterized by high point density but sparse behavior information. As shown in Figure 1(a), the aircraft usually fly in a straight line with a fixed attitude. The representative activities, such as turns, holding patterns, and takeoff/landing phases, only account for a relatively small part, *e.g.* 5%, of the whole trajectory. Existing TRL methods usually treat every point equally without considering the density of behaviors, leading to inconsequential representations.

3D spatial continuity. Flights move in a three-dimensional space, exhibiting more complex spatial patterns compared to ground trajectories. As shown in Figure 1(b), the movement of flights follows the spatial proximity in 3D space, *i.e.* the motion trend of flights are constrained in the forward direction. However, previous works utilize MSE loss to reconstruct location coordinates, which does not effectively capture the spatial continuity and dependencies. Moreover, MSE loss is sensitive to outliers, which significantly affect the representation learning process.

In this paper, we propose a generic flight trajectory representation learning method FLIGHT2VEC which encourages the learned representation paying more attention to behavior parts and memorize the motion trend of flights. Specifically, a novel behavior-based patching mechanism is designed to automatically sample trajectories according to the behavior density and reduce the sequence length of model input. FLIGHT2VEC first employs a criterion to select the behavior parts of trajectories and construct patches for each behavior respectively. The generated patches are feed to the decoder-only Transformer to obtain the representations. To model the 3D spatial continuity, we propose a motion trend learning technique which predicts the moving direction of each record as a 26-classes classification task. Along with the MSE loss, FLIGHT2VEC can not only extract the information of spatial coordinates but also capture the motion trend of flights. In summary, the main contributions of this paper are summarized as follows:

- We propose an effective and efficient representation learning framework FLIGHT2VEC for flight trajectories. To the best of our knowledge, this is the first work that designs a general representation learning framework specifically tailored to the characteristics of flight trajectories.
- We introduce an behavior-adaptive patching mechanism to achieve effective trajectory representation learning while preserving the information of behavior density. Additionally, we propose a motion trend learning technique that explicitly models the spatial continuity flight trajectories without requiring additional features.
- Extensive experiments demonstrate that FLIGHT2VEC significantly enhances performance across various downstream tasks, such as anomaly detection, trajectory prediction, and flight monitoring.

Related Work

Flight Trajectory Analysis Framework. Existing research on flight trajectory data analysis always requires manual feature engineering and specialized models for each task, heavily relying on domain expertise and being difficult to extend. For instance, (Guo et al. 2022a, 2024) proposed a feature representation method based on binary encoding (BE) for trajectory prediction, utilizing Conv1D and Transformer modules to capture spatiotemporal features of trajectory points. (Fang et al. 2021) selected 10 correlation coefficient features, such as Oil Temperature (OilT) and Oil Pressure, for trajectory recognition. (Qin et al. 2022) introduced unsupervised feature engineering methods to map input data

into latent feature spaces for anomaly detection. The limitation of these methods is that they require expert knowledge and complex feature engineering to extract useful information. There are also some learning-based methods, but they do not solve the problem of uneven behavior density. They either model the entire trajectory indiscriminately (Guo et al. 2024) or only focus on specific segments such as takeoff and landing phases (Fernández et al. 2019; Memarzadeh, Matthews, and Templin 2022), unable to adaptively identify and model informative segments of the trajectory. In conclusion, although existing methods have achieved success in specific tasks, they rely on complex feature engineering highlights the need for a unified approach.

Trajectory Representation Learning. Trajectory Representation Learning (TRL) has gained significant attention in the data engineering community due to its effectiveness in enhancing various downstream tasks. Existing methods can be broadly divided into two categories: road network-based methods and grid-based methods. Road network-based methods (Fu and Lee 2020; Chen et al. 2021; Jiang et al. 2023; Qian et al. 2024) map trajectories to nodes of the road network and learn representations of nodes to generate trajectory representations. However, flight trajectories can move in three-dimensional space without the constraints of road networks, making road network-based methods inapplicable. Grid-based methods (Yao et al. 2018, 2019; Li et al. 2018; Yang et al. 2021; Jing et al. 2022) divide the geographical space into a grid of cells and map trajectories to these grids to learn representations. These methods are also difficult to apply to flight trajectories since dividing 3D space into grids will result in an exponential increase in the number of grids. This will make the trajectory data on many grids too sparse, thus the model cannot effectively learn trajectory representations. Overall, there is no existing work specifically focused on representation learning for flight trajectories and existing TRL methods are not applicable to flight trajectories.

Patch-based Transformer. In recent years, patch-based approaches for time series analysis have emerged as mainstream. Compared to point-wise processing methods, these models not only improve efficiency but also exhibit notable performance improvements, revealing the importance of enhancing modeling of local semantics through patching. For instance, PatchTST (Nie et al. 2022) segments each time series into patches and employs an Multi-layer Perceptron (MLP) to feed patch embeddings into a Transformer. This method also utilizes mask-based unsupervised pre-training to learn representations that can generalize to various downstream tasks. Following PatchTST, a series of patch-based Transformer models for time series have continually set new benchmarks in various downstream tasks, such as prediction (Wang et al. 2024; Chen et al. 2024), classification (Li, Li, and Yan 2024) and anomaly detection (Yang et al. 2023). Notably, HDMixer (Huang et al. 2024) has revealing the importance of patch division. This research shows that incorrect patch boundaries can obscure local patterns and disrupt the semantic continuity of the sequence. However, these methods do not address the issue of unbalanced behavior

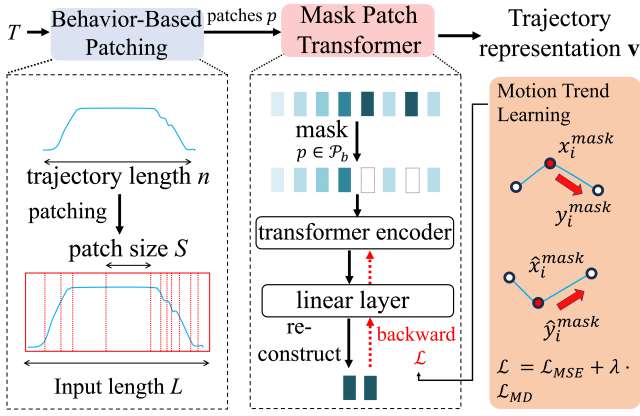


Figure 2: Overview of FLIGHT2VEC

density, *i.e.*, they fail to effectively capture sparse but critical behaviors in flight trajectories.

Preliminary

In this section, we first define the problem and then describe the proposed method FLIGHT2VEC, respectively.

Problem Definition. Given a flight trajectory denoted by $T = \{x_1, x_2, \dots, x_n\}$, where each x_i represents a recorded point at timestamp i , the objective of *Trajectory Representation Learning* (TRL) is to generate a general low-dimensional representation, $\mathbf{v} \in \mathbb{R}^D$, that can be utilized in various downstream tasks.

In this study, each trajectory point x_i is described by six key attributes related to the aircraft's status:

$$x_t = [\text{lon}_i, \text{lat}_i, \text{alt}_i, V_{\text{lon}_i}, V_{\text{lat}_i}, V_{\text{alt}_i}],$$

where lon_i , lat_i , and alt_i represent the longitude, latitude, and altitude of the aircraft, respectively. The attributes V_{lon_i} , V_{lat_i} , and V_{alt_i} denote the velocity components in the longitudinal, latitudinal, and altitude dimensions, respectively.

Overview of FLIGHT2VEC. As shown in Figure 2, FLIGHT2VEC consists of two key components, *i.e.* behavior-adaptive patching Transformer and model optimization. In the first component, we segment the flight trajectory into a sequence of patches according to the density of behaviors. For non-behavior segments, we down-sample the original records and compress the long segments into patches with equal size. The generated patches are feed to a patch Transformer encoder to obtain the flight representation. To optimize the parameters in FLIGHT2VEC, a motion trend learning approach is proposed along with the traditional MSE loss to reconstruct the masked patches. It encourages the learned representations memorize not only the location coordinates but also the moving directions of each records in the masked patches.

Methodology

We specify the two key components of FLIGHT2VEC, *i.e.*, the behavior-adaptive patching Transformer, and the optimization of the model, respectively.

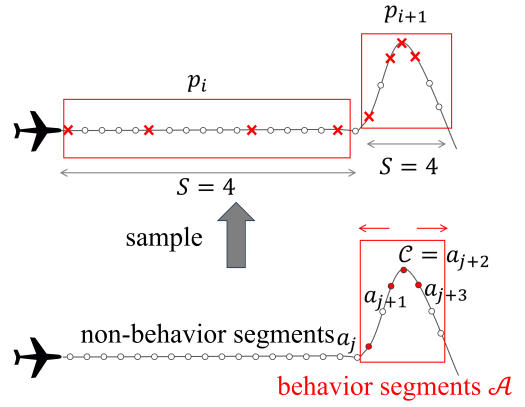


Figure 3: Behavior-Based Patching

Behavior-Adaptive Patching Transformer

Behaviors such as takeoff and turning only account for a very small part of flight trajectories, but they are informative and crucial for modeling trajectory patterns. Therefore, we propose an behavior-based patching mechanism that extracts more informative features by amplifying behavior-dense segments.

Behavior-Based Patching We observe that: (i) behavior segments reflect significant trajectory changes, such as turns and holding patterns; (ii) non-behavioral segments show stable movement with dense point distributions, which are less informative. Inspired by the observation, we first identify behavior segments \mathcal{A} in a flight trajectory T and then adaptively patch the trajectory based on \mathcal{A} .

Specifically, considering a flight trajectory indicated by $T = \{x_1, x_2, \dots, x_n\}$, we first apply a threshold filter to remove noisy points with excessive oscillations. Next, we calculate the angle change for each point as: $\text{angle}_i = \{a_1, a_2, \dots, a_n\}$.

Then, points with angle changes exceeding a threshold s are identified as active points, denoted by \mathcal{A}' :

$$\mathcal{A}' = \{a_j \mid \text{angle}_j > s\}$$

After that, we use a behavior-based patching algorithm to create patches as the inputs of Transformer. Since each behavior is composed of multiple consecutive points, for each active point, we compute its index distance to neighboring active points and cluster those with index distance less than a threshold. Then, we select a central active point c from each cluster as a patch center and we can get a patch with a predefined patch size S . As illustrated in Figure 3, a_j , a_{j+1} , and a_{j+2} are all active points. Since they are adjacent, we consider them to be in the same cluster and select point a_{j+1} as the center point c . After that, we obtain a behavioral patch set \mathbf{P}_b consists of g patches, where g is the number of active point clusters. Each patch in \mathbf{P}_b preserves the information of a behavioral segment, improving the model's ability to learn from local patterns.

For non-behavioral segments, we perform uniform sampling with a step size of $\frac{n-S \cdot g}{N-g}$ to generate additional patches




Figure 4: Illustration of Motion Trend Learning

from these segments, where the N denotes the number of patches and S is the patch size. Overall, we get a sequence of patches $\mathbf{P} = [p_1, \dots, p_N]$ which consists of g behavioral patches and $N - g$ non-behavioral patches.

Patch Transformer FLIGHT2VEC utilizes a similar approach similar to PatchTST (Nie et al. 2022) to learn the trajectory embeddings.

Specifically, we transform each patch p_i into the latent space of dimension d_m via a trainable linear projection \mathcal{W}_p , and add a learnable additive position encoding matrix \mathcal{W}_{pos} to encode its temporal order:

$$\mathbf{Y} = \mathbf{W}_p \mathbf{P} + \mathbf{W}_{pos}$$

Then, we feed \mathbf{Y} into a Transformer to generate the final trajectory embedding $\mathbf{v} = \mathbf{Transformer}(\mathbf{Y})$, where $\mathbf{v} \in \mathbb{R}^{N \times S \times d_m}$.

Optimization of FLIGHT2VEC

To model the complex moving patterns and the spatial continuity in flight trajectories, we optimize FLIGHT2VEC with a mask-based self-supervised learning strategy and a novel motion trend learning.

Mask-based self-supervised learning. Previous methods often adopt random masking to optimize patch Transformer, which is not suitable for flight trajectories where most points are moving uniformly in a straight line. For these points, the model can easily infer the masked values through simple interpolation between adjacent time points, without capturing the complex trajectory patterns. To address this problem, we introduce a motion-based randomization strategy to effectively mask the behavior patches \mathbf{P}_b and their surrounding areas. Specifically, we mask the patches in \mathbf{P}_b and their neighboring patches with the probability ρ_b , and we mask other patches with the probability ρ_n , where $\rho_n < \rho_b$. Then, we use a $D \times P$ linear layer to reconstruct the masked patches x^{mask} . The Mean Squared Error (MSE) loss to minimize reconstruction errors:

$$\mathcal{L}_{MSE} = \mathbb{E}_x \frac{1}{M} \sum_{i=1}^M \|\hat{x}_i^{masked} - x_i^{mask}\|^2$$

where M is the number of points in the masked patches.

Motion Trend Learning. However, MSE loss is not enough to model flight trajectories for two main reasons. First, MSE loss treats all errors equally, regardless of their spatial context, which means it does not effectively capture

the spatial continuity and dependencies inherent in flight trajectories. Second, MSE loss does not adequately model the uncertainty and sparsity of significant behavioral points within dense trajectory data. Critical behavioral segments, such as turns and altitude changes, are sparse but vital for accurate representation. Optimizing the model with only MSE loss may result in suboptimal learning, as it might not adequately focus on these sparse yet crucial points.

For example, as illustrated in Figure 4, consider a point x_1 that needs to be reconstructed or predicted, with two candidate points x_a and x_b . Both x_a and x_b are equidistant from x_1 , resulting in the same MSE loss. But x_a is better than x_b because it is located on the line where the trajectory is moving forward. To model this preference and inherent uncertainty, we introduce the moving direction loss function.

Specifically, we categorize each point in the trajectory based on its movement direction. For point $x_i = (lon_i, lat_i, alt_i)$ and $x_{i+1} = (lon_{i+1}, lat_{i+1}, alt_{i+1})$, we calculate the direction vector \vec{d} from x_i to x_{i+1} :

$$\vec{d} = (lon_{i+1} - lon_i, lat_{i+1} - lat_i, alt_{i+1} - alt_i).$$

Each component of \vec{d} can be positive, negative, or zero, corresponding to the direction along each axis. The direction is represented as a triplet $(d_{lon}, d_{lat}, d_{alt})$ based on the direction vector, where $d_{lon} = \text{sign}(lon_{i+1} - lon_i)$, $d_{lat} = \text{sign}(lat_{i+1} - lat_i)$, and $d_{alt} = \text{sign}(alt_{i+1} - alt_i)$. The sign function is defined as:

$$\text{sign}(u) = \begin{cases} 1 & \text{if } u > 0 \\ 0 & \text{if } u = 0 \\ -1 & \text{if } u < 0 \end{cases}$$

We divided the 3D moving space into 26 categories based on directional movement. Each dimension (longitude, latitude, and altitude) has three possible states: positive, negative, or unchanged. Combining these states yields $3 \times 3 \times 3 - 1 = 26$ unique directions, excluding the case where all dimensions remain unchanged. Each point is then assigned a category corresponding to its movement direction.

We adopt following moving direction loss to optimize the model and to learn these directional preferences:

$$\mathcal{L}_{MD} = -\mathbb{E}_x \log \mathbb{P}(y_i^{masked} | y_{1:i-1}^{mask}, x)$$

where y_i^{masked} represents the label of one of the 26 moving directions.

The final loss is a combination of Mean Squared Error (MSE) loss and the moving direction loss, enabling the model to capture both spatial relationships and precise values. Thus, the final combined loss function is:

$$\mathcal{L}_m = \mathcal{L}_{MD} + \lambda \cdot \mathcal{L}_{MSE}$$

where λ is a weighting factor that balances the contributions of the two parts.

By incorporating the spatial proximity-aware loss function and the moving direction loss, FLIGHT2VEC can effectively encode the complex three-dimensional movement patterns in flight trajectories into the learned representations.

Methods	Hor.	MAE ↓			MAPE (%) ↓			RMSE ↓			MDE ↓
		Lon	Lat	Alt	Lon	Lat	Alt	Lon	Lat	Alt	
PatchTST	1	0.0041	0.0039	1.47	0.0032	0.0136	0.32	0.0125	0.0134	8.04	0.79
	3	0.0073	0.0072	2.98	0.0063	0.0256	0.61	0.0214	0.0231	8.72	1.34
	15	0.0126	0.0129	8.03	0.0125	0.0532	1.77	0.0305	0.0322	19.94	1.61
	30	0.0231	0.0248	14.99	0.0217	0.0962	3.48	0.0591	0.0672	38.88	2.70
	60	0.0555	0.0572	24.65	0.0489	0.1978	7.87	0.1129	0.1376	72.15	6.50
FlightBERT++	1	0.0018	0.0018	1.16	0.0016	0.0067	0.21	0.0037	0.0116	13.11	0.32
	3	0.0032	0.0032	2.33	0.0031	0.0119	0.43	0.0076	0.0133	12.88	0.59
	15	0.0125	0.0118	7.48	0.0110	0.0429	1.38	0.0268	0.0329	22.91	1.99
	30	0.0214	0.0213	12.97	0.0286	0.0916	2.46	0.0576	0.0671	48.55	4.44
	60	0.0482	0.0497	29.87	0.0606	0.1666	6.02	0.1076	0.1319	79.25	7.24
FLIGHT2VEC	1	0.0019	0.0018	1.17	0.0015	0.0062	0.23	0.0038	0.0121	8.02	0.34
	3	0.0034	0.0033	2.32	0.0029	0.0118	0.42	0.0075	0.0136	8.11	0.57
	15	0.0124	0.0121	7.58	0.0109	0.0431	1.36	0.0259	0.0319	18.24	1.98
	30	0.0196	0.0213	12.14	0.0212	0.0811	2.09	0.0436	0.0501	37.52	3.34
	60	0.0381	0.0401	25.67	0.0403	0.1216	3.72	0.0836	0.1019	56.45	5.14
FLIGHT2VEC +BE	1	0.0017	0.0017	1.15	0.0015	0.0061	0.19	0.0036	0.0112	8.01	0.29
	3	0.0029	0.0029	1.92	0.0028	0.0116	0.39	0.0036	0.0109	7.83	0.55
	15	0.0122	0.0112	7.43	0.0101	0.0402	1.31	0.0243	0.0312	19.81	1.44
	30	0.0192	0.0198	11.89	0.0203	0.0699	2.03	0.0432	0.0498	33.88	2.24
	60	0.0371	0.0392	25.45	0.0401	0.1209	3.62	0.0819	0.0932	54.77	5.04

Table 1: The experimental results of FTP task.

Complexity Analysis

The complexity of the behavior-based patching is $O(n)$, where n represents the length of the flight trajectory. Then, the complexity of patch projection and position embedding is $O(N \cdot S \cdot d_m)$, where N, S, d_m denote the number of patches, the size of each patch and the embedding dimension, respectively. Then, the backbone, *i.e.*, the Transformer encoder, involves self-attention computation, resulting in a complexity of $O(N^2 \cdot d_m)$. Finally, the complexity of the Linear Layer is $O(N \cdot d_m)$. In general, the time complexity of our framework is $O(N^2 \cdot d_m)$, with the Transformer Encoder being the most computationally intensive component. In practice, the number of patches is much less the length of trajectory which makes FLIGHT2VEC efficient.

Experiment

Experimental Settings

In this section, we briefly introduce the datasets, experiment protocols, baselines and hyperparameter settings.

Data Descriptions We conduct extensive experiments on two real-world datasets, the Swedish Civil Air Traffic Control (SCAT) (Nilsson and Unger 2023) and Aircraft Trajectory Classification Data for Air Traffic Management(ATFMTraj) (Phisannupawong, Damanik, and Choi 2024b).

Experimental Protocol We employ three representative tasks, *i.e.* flight trajectory prediction (FTP), flight recognition (FR), and anomaly detection (AD), to evaluate the representations learned from FLIGHT2VEC. For FTP, we pre-

dict the future trajectory in (1, 3, 15, 30, 60) different horizons and utilize Mean Absolute Error (MAE), Mean Absolute Percentage Error (MAPE), Root Mean Squared Error (RMSE) and Mean Distance Error (MDE) as the evaluation metrics. For FR, we directly use the category of flight as ground truth and evaluate the performance with accuracy (ACC), precision (PRE), and recall (REC). For AD, we generate the synthetic anomalies according to the previous work (Guo et al. 2022b) and evaluate the performance of FLIGHT2VEC with AUC and AUPR. Following the settings of the previous work, we use the Mean Time Cost(MTC) metric to evaluate the computational performance.

Baselines As described in the experimental protocol, we employ three tasks to verify the performance of FLIGHT2VEC. For flight trajectory prediction, two representative methods, *i.e.* **FlightBERT++** (Guo et al. 2024) and **PatchTST** (Nie et al. 2022) are compared. For anomaly detection methods, we use **DMDN** (Lijing, Weili, and Zhao 2021), and **DDM** (Guo et al. 2022b) for performance comparison. For flight trajectory recognition, we select **SPIRAL** (Lei et al. 2019) and **ATSCC** (Phisannupawong, Damanik, and Choi 2024a) as our baselines. Moreover, we also compare the ablations of FLIGHT2VEC to verify the superiority of the proposed techniques and study the parameter sensitivity to provide some insight in the use of FLIGHT2VEC.

Hyperparameters setting Our model utilizes the Transformer configuration from (Nie et al. 2022), which includes 3 layers with a model dimension of 256, and 16 attention heads with a dropout rate of 0.2. The binomial masking

Method	AUC				AUPR			
	SMA	HD	VD	Go-Around	SMA	HD	VD	Go-Around
DMDN	0.7498	0.7444	0.7354	0.7100	0.7416	0.7498	0.7326	0.7210
DDM	0.8509	0.8411	0.8324	0.8200	0.8122	0.8112	0.7923	0.7890
FLIGHT2VEC	0.9201	0.9178	0.9127	0.9050	0.9072	0.9059	0.9017	0.9005

Table 2: The experimental results of anomaly detection task.

Methods	RKSla			RKSId			ESSA			LSZH		
	ACC	PRE	REC	ACC	PRE	REC	ACC	PRE	REC	ACC	PRE	REC
SPIRAL	0.8344	0.8301	0.8302	0.9946	0.9942	0.9932	0.8503	0.8496	0.8408	0.9173	0.9172	0.9124
ATSCC	0.9946	0.9951	0.9911	0.9987	0.9981	0.9971	0.9990	0.9989	0.9980	0.9977	0.9979	0.9971
Flight2Vec	0.9947	0.9952	0.9943	0.9988	0.9980	0.9970	0.9990	0.9986	0.9982	0.9981	0.9980	0.9980

Table 3: The experimental results of flight recognition task.

probability is set at 0.4. The dimension of the representation p_i is set to 256. For training, the batch size is set to 256, and the AdamW optimizer is used with a learning rate of 1×10^{-5} . The model is pre-trained for 100 epochs with a patch length of 32. All the experiments are conducted on the $2 \times$ NVIDIA 3090Ti.

Effectiveness of FLIGHT2VEC

To evaluate the effectiveness of FLIGHT2VEC, we conduct experiments on the aforementioned three tasks and analyze the results of each task respectively.

Results of Flight Trajectory Prediction. We conduct FTP experiment on SCAT dataset and report the quantitative results in Table 1. According to the results, we have three observations. Firstly, FlightBERT++ is the most competitive baseline, which can predict the flight trajectory in a non-autoregressive manner, but it is also inferior to FLIGHT2VEC. This proves the effectiveness of activity density and moving trends modeled in FLIGHT2VEC. Secondly, with the prediction horizon increasing, the performances of all methods are dropped. For example, the MDE of FlightBERT++ increases by approximately 237% when the horizon increases from 3 to 15. FLIGHT2VEC still beats all compared baselines on all prediction horizons. We attribute this to the effectiveness of moving direction loss which captures the 3D spatial motion trend of flight trajectory. Lastly, due to the fusion of some attributes of the trajectory points, FlightBERT++ achieves better performance than FLIGHT2VEC for short-term horizon predictions (in horizons 1 and 3). We use the same method in FlightBERT++ to integrate the attributes in FLIGHT2VEC and form FLIGHT2VEC +BE. As shown, FLIGHT2VEC +BE achieves the best performance across all prediction horizons, which also proves the scalability of our method.

Results of Anomaly Detection. The flight trajectory anomaly detection experiments are also conducted on SCAT dataset. The results are presented in Table 2. Due to the absence of real ground truth anomalies. We construct four types of anomalies, *i.e.*, Successive Multipoint Anomaly (SMA), Horizontal Deviation (HD), Vertical Deviation (VD)

and Go-Around following the existing work (Guo et al. 2022b). We train another MLP layer along with the learned representations of FLIGHT2VEC to achieve anomaly detection. All the methods are trained with the same training dataset. For evaluation metrics, we utilize AUPR in addition to AUC to handle the unbalanced anomaly classes. As shown, the performance of DMDN is inferior to DDM, which proves the density estimation method of DMDN can not learn representative features for detecting anomalies. FLIGHT2VEC achieves significant improvements in all metrics, proving that the learned representations can capture the local movement patterns of the flight trajectory. For Go-Around anomaly, FLIGHT2VEC achieves 10.4% and 14.1% improvements of AUC and AUPR, compared with the most competitive baseline DDM. Moreover, the results of FLIGHT2VEC consistently outperform the baselines on four anomaly types, indicating that the proposed behavior-adaptive patching and moving direction loss can encourage the representation model to learn useful features.

Results of Flight Recognition. We conduct the flight recognition experiment on ATFMTraj dataset and present the results in Table 3. In ATFMTraj, the flight trajectories are categorized into four classes, *i.e.*, RKSla, RKSId, ESSA and LSZH, by the airports of take-off and landing. As illustrated, the performances of all methods are relatively high. The best accuracy of the four classes are over 99%. For the compared baselines, the performance of ATSCC significantly outperforms SPIRAL. The precision increased from 0.8301 to 0.9946 on RKSla. This phenomenon indicates the segment-patch framework in ATSCC can better represent the flight trajectory compared with using original data directly. FLIGHT2VEC utilizes the behavior-adaptive patching Transformer and achieves comparable results with the performance of ATSCC.

Efficiency of FLIGHT2VEC

To verify the efficiency of FLIGHT2VEC, we report the Mean Time Cost (MTC) of representation generation along with the model size in Table 4. As shown, the computational time of FLIGHT2VEC is better than, at least com-

Task	Methods	Parameters (M)	MTC (ms)
FTP	LSTM+Attention	0.90	48.96
	PatchTST	1.61	2.20
	FlightBERT++	29.96	7.21
	FLIGHT2VEC	1.86	3.12
FR	SPIRAL	-	0.76
	ATSCC	0.90	43.89
	FLIGHT2VEC	1.85	3.11
AD	DMDN	0.61	47.96
	DDM	0.11	6.53
	FLIGHT2VEC	1.85	3.11

Table 4: Comparison of the Computational Performance.

Methods	Hor.	MAE ↓			MAPE (%) ↓			RMSE ↓		
		Lon	Lat	Alt	Lon	Lat	Alt	Lon	Lat	Alt
w/o PD	1	0.0032	0.0032	1.65	0.0029	0.0098	0.54	0.0052	0.0145	14.55
	15	0.0119	0.0117	7.94	0.0117	0.0512	1.62	0.0289	0.0276	19.12
w/o MD	1	0.0022	0.0022	1.39	0.0019	0.0072	0.29	0.0042	0.0133	13.95
	15	0.0144	0.0141	7.88	0.0162	0.0438	1.42	0.0279	0.0382	23.11
our	1	0.0019	0.0018	1.17	0.0015	0.0062	0.23	0.0038	0.0121	13.21
	15	0.0124	0.0121	7.58	0.0109	0.0431	1.36	0.0259	0.0319	18.24

Table 5: The experimental results of the ablation study.

parable with, all the compared baselines. For the FTP task, FLIGHT2VEC is the fast method expect for PatchTST. However, PatchTST is a light model that is not specifically designed for flight trajectories and performs poorer than FLIGHT2VEC. FlightBERT++ is inferior to FLIGHT2VEC, because FlightBERT++ is a encoder-decoder framework while FLIGHT2VEC is a decoder-only architecture. The parameters in FLIGHT2VEC are much less than FlightBERT++, leading to better computational efficiency. For FR, SPIRAL has minimal computational time. This is because SPIRAL is not a deep learning method. Compared with the SOTA method, FLIGHT2VEC is 10 times faster than ATSCC, *i.e.*, from 43.89 ms (ATSCC) to 3.11 ms (FLIGHT2VEC). For AD task, FLIGHT2VEC demonstrate substantial improvements in computational performance compared to the baselines.

Ablation Study

We compare FLIGHT2VEC with two ablations to analyze the effectiveness of the proposed components. We remove the proposed behavior-based patching, randomly sample points at fixed intervals and divide the patch to obtain w/o PD. We obtain w/o MD by removing moving direction loss.

Due to the space limit, we only report the experiment on FTP task and the results are shown in Table 5. We observe: (1) Comparing the results of FLIGHT2VEC with w/o MD, we observe the moving direction loss can model the spatial continuity flight trajectories. For example, the RMSE improves from 18.24 to 23.11 on altitude. (2) From the results of w/o PD and FLIGHT2VEC, we can conclude that the ac-

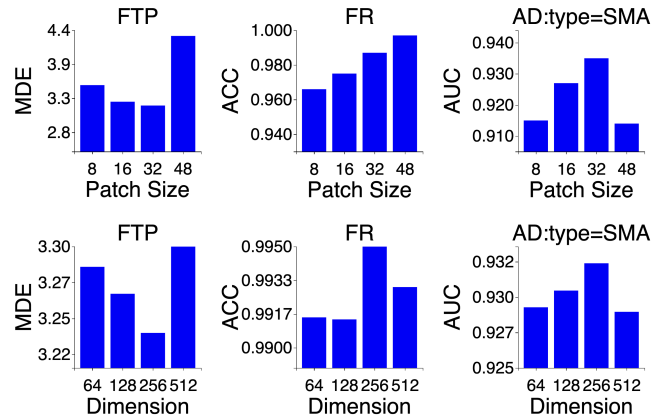


Figure 5: MDE scores with varying model parameters.

tivity patching mechanism capacity to represent flight trajectory is more effective than a fixed patch. (3) FLIGHT2VEC achieves the best performance compared to all ablations, which proves the effectiveness of the proposed techniques.

Sensitivity Analysis

We analyze the impact of selecting different patch sizes and representation dimensions on different tasks. We ranged the patch size from 8 to 48 and the representation dimension from 64 to 512, and then show the average performance of these configurations on different tasks in Figure 5.

With the increase of patch size, the performance of FLIGHT2VEC first increases and then drops. FLIGHT2VEC achieves the best performance with the patch size of 32 on both FTP and AD tasks. For FR task, the optimal performance is achieved with the largest patch size of 48. This is because the FR task pays more attention to the global movement patterns of the flight trajectory.

For the change of embedding dimension, the performances of the three tasks have similar patterns. FLIGHT2VEC achieves the best performance on the dimension size of 256. This phenomenon indicates FLIGHT2VEC is robust to different embedding dimension sizes.

Conclusion

In this paper, we present FLIGHT2VEC which is the first unified framework specifically designed for flight trajectory representation learning. By addressing the challenges of unbalanced behavior density and 3D spatial continuity, FLIGHT2VEC enhances the representation of flight trajectories. Our approach introduces an adaptive behavior-based patching mechanism and a direction loss, significantly improving performance across tasks like trajectory prediction, anomaly detection, and flight monitoring. Extensive experiments across multiple task demonstrate that FLIGHT2VEC not only significantly outperforms existing methods but also sets a new benchmark in the field.

Acknowledgments

This work has been supported by the National Natural Science Foundation of China under Grant No. 62472405. This work is also sponsored by CCF-DiDi GAIA Collaborative Research Funds for Young Scholars.

References

- Chen, P.; Zhang, Y.; Cheng, Y.; Shu, Y.; Wang, Y.; Wen, Q.; Yang, B.; and Guo, C. 2024. Multi-scale transformers with adaptive pathways for time series forecasting. In *International Conference on Learning Representations*.
- Chen, Y.; Li, X.; Cong, G.; Bao, Z.; Long, C.; Liu, Y.; Chandran, A. K.; and Ellison, R. 2021. Robust road network representation learning: When traffic patterns meet traveling semantics. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 211–220.
- Fang, W.; Wang, Y.; Yan, W.; and Lin, C. 2021. Symbolic flight action recognition based on neural networks. *Syst. Eng. Electron.*, 13: 963–969.
- Fernández, A.; Martínez, D.; Hernández, P.; Cristóbal, S.; Schwaiger, F.; Nunez, J. M.; and Ruiz, J. M. 2019. Flight data monitoring (FDM) unknown hazards detection during approach phase using clustering techniques and AutoEncoders. *Proceedings of the Ninth SESAR Innovation Days, Athens, Greece*, 2–5.
- Fu, T.-Y.; and Lee, W.-C. 2020. Trembr: Exploring road networks for trajectory representation learning. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(1): 1–25.
- Guo, D.; Wu, E. Q.; Wu, Y.; Zhang, J.; Law, R.; and Lin, Y. 2022a. FlightBERT: binary encoding representation for flight trajectory prediction. *IEEE Transactions on Intelligent Transportation Systems*, 24(2): 1828–1842.
- Guo, D.; Zhang, Z.; Yan, Z.; Zhang, J.; and Lin, Y. 2024. FlightBERT++: A Non-autoregressive Multi-Horizon Flight Trajectory Prediction Framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 127–134.
- Guo, Z.; Yin, C.; Zeng, W.; Tan, X.; and Bao, J. 2022b. Data-driven method for detecting flight trajectory deviation anomaly. *Journal of Aerospace Information Systems*, 19(12): 799–810.
- Huang, Q.; Shen, L.; Zhang, R.; Cheng, J.; Ding, S.; Zhou, Z.; and Wang, Y. 2024. Hdmixer: Hierarchical dependency with extendable patch for multivariate time series forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12608–12616.
- Jiang, J.; Pan, D.; Ren, H.; Jiang, X.; Li, C.; and Wang, J. 2023. Self-supervised trajectory representation learning with temporal regularities and travel semantics. In *2023 IEEE 39th international conference on data engineering (ICDE)*, 843–855. IEEE.
- Jing, Q.; Liu, S.; Fan, X.; Li, J.; Yao, D.; Wang, B.; and Bi, J. 2022. Can Adversarial Training benefit Trajectory Representation? An Investigation on Robustness for Trajectory Similarity Computation. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, 905–914.
- Lei, Q.; Yi, J.; Vaculin, R.; Wu, L.; and Dhillon, I. S. 2019. Similarity preserving representation learning for time series clustering. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 2845–2851.
- Li, X.; Zhao, K.; Cong, G.; Jensen, C. S.; and Wei, W. 2018. Deep representation learning for trajectory similarity computation. In *2018 IEEE 34th international conference on data engineering (ICDE)*, 617–628. IEEE.
- Li, Z.; Li, S.; and Yan, X. 2024. Time series as images: Vision transformer for irregularly sampled time series. *Advances in Neural Information Processing Systems*, 36.
- Lijing, C.; Weili, Z.; and Zhao, Y. 2021. An Aircraft Trajectory Anomaly Detection Method Based on Deep Mixture Density Network. *Transactions of Nanjing University of Aeronautics & Astronautics*, 38(5).
- Liu, Y.; and Hansen, M. 2018. Predicting aircraft trajectories: A deep generative convolutional recurrent neural networks approach. *arXiv preprint arXiv:1812.11670*.
- Memarzadeh, M.; Matthews, B.; and Templin, T. 2022. Multiclass Anomaly Detection in Flight Data Using Semi-Supervised Explainable Deep Learning Model. *Journal of Aerospace Information Systems*, 19(2): 83–97.
- Memarzadeh, M.; Matthews, B. L.; and Weckler, D. I. 2023. Anomaly Detection in Flight Operational Data Using Deep Learning. In *System-Wide Safety Technical Challenge 1 Close Out Event*.
- Nie, Y.; Nguyen, N. H.; Sinthong, P.; and Kalagnanam, J. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730*.
- Nilsson, J.; and Unger, J. 2023. Swedish civil air traffic control dataset. *Data in brief*, 48: 109240.
- Olive, X.; and Basora, L. 2020. Detection and identification of significant events in historical aircraft trajectory data. *Transportation Research Part C: Emerging Technologies*, 119: 102737.
- Phisannupawong, T.; Damanik, J. J.; and Choi, H.-L. 2024a. Aircraft Trajectory Segmentation-based Contrastive Coding: A Framework for Self-supervised Trajectory Representation. *arXiv:2407.20028*.
- Phisannupawong, T.; Damanik, J. J.; and Choi, H.-L. 2024b. ATFMTraj: Aircraft Trajectory Classification Data for Air Traffic Management. <https://huggingface.co/datasets/petchthwr/ATFMTraj>.
- Qian, T.; Li, J.; Chen, Y.; Cong, G.; Sun, T.; Wang, F.; and Xu, Y. 2024. Context-Enhanced Multi-View Trajectory Representation Learning: Bridging the Gap through Self-Supervised Models. *arXiv preprint arXiv:2410.13196*.
- Qin, K.; Wang, Q.; Lu, B.; Sun, H.; and Shu, P. 2022. Flight anomaly detection via a deep hybrid model. *Aerospace*, 9(6): 329.
- Wang, S.; Wu, H.; Shi, X.; Hu, T.; Luo, H.; Ma, L.; Zhang, J. Y.; and Zhou, J. 2024. Timemixer: Decomposable multiscale mixing for time series forecasting. *arXiv preprint arXiv:2405.14616*.

- Wu, X.; Yang, H.; Chen, H.; Hu, Q.; and Hu, H. 2022. Long-term 4D trajectory prediction using generative adversarial networks. *Transportation Research Part C: Emerging Technologies*, 136: 103554.
- Yang, P.; Wang, H.; Zhang, Y.; Qin, L.; Zhang, W.; and Lin, X. 2021. T3s: Effective representation learning for trajectory similarity computation. In *2021 IEEE 37th International Conference on Data Engineering (ICDE)*, 2183–2188. IEEE.
- Yang, Y.; Zhang, C.; Zhou, T.; Wen, Q.; and Sun, L. 2023. Dcdetector: Dual attention contrastive representation learning for time series anomaly detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 3033–3045.
- Yao, D.; Cong, G.; Zhang, C.; and Bi, J. 2019. Computing trajectory similarity in linear time: A generic seed-guided neural metric learning approach. In *2019 IEEE 35th international conference on data engineering (ICDE)*, 1358–1369. IEEE.
- Yao, D.; Hu, H.; Du, L.; Cong, G.; Han, S.; and Bi, J. 2022. TrajGAT: A graph-based long-term dependency modeling approach for trajectory similarity computation. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, 2275–2285.
- Yao, D.; Zhang, C.; Zhu, Z.; Hu, Q.; Wang, Z.; Huang, J.; and Bi, J. 2018. Learning deep representation for trajectory clustering. *Expert Systems*, 35(2): e12252.
- Yao, D.; Zhang, C.; Zhu, Z.; Huang, J.; and Bi, J. 2017. Trajectory clustering via deep representation learning. In *2017 international joint conference on neural networks (IJCNN)*, 3880–3887. IEEE.
- Zhang, X.; Yin, Z.; Liu, F.; and Huang, Q. 2016. Data mining method for aircraft maneuvering division. *J. Northwest. Polytech. Univ.*, 34: 33–40.
- Zhang, Z.; Guo, D.; Zhou, S.; Zhang, J.; and Lin, Y. 2023. Flight trajectory prediction enabled by time-frequency wavelet transform. *Nature Communications*, 14(1): 5258.