

# X-VORTEX: Spatio-Temporal Contrastive Learning for Wake Vortex Trajectory Forecasting

Zhan Qu  
TU Dresden, ScaDS.AI  
Dresden, Germany  
zhan.qu@tu-dresden.de

Michael Färber  
TU Dresden, ScaDS.AI  
Dresden, Germany  
michael.farber@tu-dresden.de

## Abstract

Wake vortices are strong, coherent air turbulences created by aircraft, and they pose a major safety and capacity challenge for air traffic management. Tracking how vortices move, weaken, and dissipate over time from LiDAR measurements is still difficult because scans are sparse, vortex signatures fade as the flow breaks down under atmospheric turbulence and instabilities, and point-wise annotation is prohibitively expensive. Existing approaches largely treat each scan as an independent, fully supervised segmentation problem, which overlooks temporal structure and does not scale to the vast unlabeled archives collected in practice. We present X-VORTEX, a spatio-temporal contrastive learning framework grounded in Augmentation Overlap Theory that learns physics-aware representations from unlabeled LiDAR point cloud sequences. X-VORTEX addresses two core challenges: sensor sparsity and time-varying vortex dynamics. It constructs paired inputs from the same underlying flight event by combining a weakly perturbed sequence with a strongly augmented counterpart produced via temporal subsampling and spatial masking, encouraging the model to align representations across missing frames and partial observations. Architecturally, a time-distributed geometric encoder extracts per-scan features and a sequential aggregator models the evolving vortex state across variable-length sequences. We evaluate on a real-world dataset of over one million LiDAR scans. X-VORTEX achieves superior vortex center localization while using only 1% of the labeled data required by supervised baselines, and the learned representations support accurate trajectory forecasting.

## CCS Concepts

• Applied computing → Aerospace; • Computing methodologies → Image segmentation.

## Keywords

Self-Supervised Learning; LiDAR Point Clouds; Wake Vortex; Contrastive Learning

### ACM Reference Format:

Zhan Qu and Michael Färber. 2026. X-VORTEX: Spatio-Temporal Contrastive Learning for Wake Vortex Trajectory Forecasting. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3770855.3818507>



This work is licensed under a Creative Commons Attribution 4.0 International License. *KDD '26, Jeju Island, Republic of Korea*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2259-2/2026/08  
<https://doi.org/10.1145/3770855.3818507>


## 1 Introduction

Aircraft wake vortices are coherent, counter-rotating airflows generated at the wingtips of lifting aircraft (see Fig. 1). These vortices can persist for tens of seconds and drift several hundred meters behind the generating aircraft, posing a critical safety hazard for following traffic during takeoff and landing [14, 16]. The danger posed by these encounters is not merely theoretical; it has led to catastrophic loss of life and severe operational disruptions.

On November 12, 2001, American Airlines Flight 587 crashed shortly after takeoff from JFK Airport, killing all 265 people on board and 5 on the ground. The National Transportation Safety Board (NTSB) determined that the accident was precipitated by a wake vortex encounter with a preceding Boeing 747, which induced loads that snapped the Airbus A300's vertical stabilizer in mid-flight [25]. More recently, on January 7, 2017, a Bombardier Challenger 604 business jet operating over the Arabian Sea was caught in the wake of an Airbus A380 flying 1,000 feet above. The encounter caused the jet to roll uncontrollably and plummet 9,000 feet, resulting in extreme G-forces and a total hull loss [5].

To mitigate such risks, international regulations mandate strict minimum separation distances between aircraft [9]. While effective at preventing accidents, these static separations are conservative and significantly constrain airport throughput [20]. This creates a critical trade-off between operational safety and airspace capacity. Consequently, reliable wake vortex detection and precise trajectory forecasting are essential to safely reduce separations and optimize runway efficiency, a capability that X-VORTEX aims to provide.

Recent advances in scanning LiDAR systems have enabled direct observation of wake vortices in real-world environments [21]. Each LiDAR sweep produces a 3D point cloud of the radial velocity field (see Fig. 1), and consecutive sweeps naturally form time series that



**Figure 1: Measurement geometry of a ground-based LiDAR (L) scanning perpendicular to the runway. The scan captures the radial velocity of the Port and Starboard vortices trailing a landing aircraft [33, 43].**


capture the evolution of wake vortex trajectories. However, existing learning-based approaches analyze these scans independently, treating each sweep as a static snapshot [33, 38, 43]. By relying on per-frame segmentation or detection models, prior methods largely ignore the temporal dynamics of wake evolution, such as continuous descent, decay, and dissipation driven by fluid instabilities. Moreover, supervised training of these models requires dense point-wise annotations for individual frames, which are prohibitively expensive to generate at scale [35].

Modeling these temporal dynamics is vital for three reasons. First, wake vortices evolve continuously through processes such as descent, deformation, and decay, so a vortex’s current geometry encodes rich information about its underlying dynamics, which is discarded by frame-wise analysis. Second, the vast majority of collected LiDAR data remains unlabeled. Leveraging temporal consistency provides a natural self-supervisory signal that can unlock the value of millions of archived scans [20]. Third, operational wake advisory systems require more than instantaneous detection. They depend on short-term trajectory forecasting to anticipate hazardous regions and safely optimize aircraft separation, making future-state prediction central to practical deployment.

To address these challenges, we introduce **X-VORTEX** (see Fig. 3 for an overview), a spatio-temporal contrastive learning framework for wake vortex analysis. Unlike previous supervised approaches, X-VORTEX is grounded in the **Augmentation Overlap Theory** [42, 46], which posits that semantically consistent augmentations induce overlapping views of the same underlying event, allowing contrastive learning to connect diverse observations and learn robust representations. We operationalize this by treating "temporal evolution" as a natural augmentation: we train the model to recognize that a spatially masked scan at time  $t$  and a temporally subsampled sequence at time  $t + \Delta t$  represent the same underlying physical event. X-VORTEX employs a Time-Distributed geometric encoder (e.g., PointNet) coupled with a sequential module (e.g., LSTM) to extract a physics-aware global representation from unlabeled sequences. By optimizing a multi-view InfoNCE objective, the model learns to filter out transient sensor noise while preserving the stable geometric features of the vortex cores. This pre-trained representation is then adapted to two downstream tasks: robust center localization (using a differentiable soft-center mechanism) and future trajectory forecasting.

We evaluate X-VORTEX on a massive dataset of over one million LiDAR scans collected at Vienna International Airport (LOWW) [20]. Our experiments demonstrate that by leveraging large-scale unlabeled data, X-VORTEX significantly outperforms supervised baselines, particularly in low-data regimes. Moreover, it enables the first effective short-term forecasting of wake vortex trajectories from point cloud sequences. X-VORTEX is publicly available.<sup>1</sup> In summary, the contributions of this paper are as follows:

- **Contrastive Pre-training for Wake Vortices:** We introduce the first self-supervised contrastive framework for aircraft generated wake vortices, learning sequence-level representations from unlabeled Doppler LiDAR data.



**Figure 2: Large Eddy Simulation (LES) visualizing the roll-up process. Vorticity sheets shed from the wingtips and flaps merge to form the characteristic counter-rotating vortex pair [33, 43].**

- **Dynamics-Aware View Construction:** We propose a two-view augmentation strategy combining temporal subsampling and spatial masking to enforce invariance to spatial sparsity and temporal gaps while preserving the geometric dynamics of wake vortices.
- **Wake Trajectory Forecasting:** We demonstrate the first short-horizon wake vortex trajectory forecasting from LiDAR point cloud sequences, outperforming trajectory-only and kinematic baselines.
- **Improved Label Efficiency:** We show that large-scale self-supervised pre-training substantially reduces annotation requirements, enabling accurate localization with only 1% labeled data.

## 2 Background

### 2.1 LiDAR Remote Sensing for Aviation

Light Detection and Ranging (LiDAR) has emerged as the gold standard for clear-air turbulence detection. Unlike radar, which relies on Rayleigh scattering from precipitation droplets (and thus suffers in clear air), Doppler LiDAR relies on Mie scattering from microscopic aerosols. This allows for high-precision measurements of wind fields even in visibly clear conditions [23].

The instrument operates by emitting coherent laser pulses and analyzing the frequency shift of the back-scattered signal via the Doppler effect; this shift reveals the radial velocity of the air particles along the line of sight. By scanning a vertical slice perpendicular to the runway (Range-Height-Indicator scan), the LiDAR constructs a high-resolution 2D cross-section of the atmosphere (see Fig. 1). A sequential series of these scans effectively captures the spatio-temporal evolution of flow structures. This capacity for real-time, high-spatial-resolution velocity mapping makes LiDAR indispensable for Wake Vortex Advisory Systems, where precise localization of invisible turbulence is required to optimize aircraft separation standards safely.

### 2.2 Wake Vortex Physics and Ground Effects

Wake vortices are an inevitable byproduct of lift generation. As an aircraft generates lift, the pressure differential between the lower

<sup>1</sup><https://github.com/ZhanQu945/X-Vortex>

(high pressure) and upper (low pressure) wing surfaces induces a spanwise flow, causing vortex sheets to shed from the wingtips and flaps [37]. These sheets rapidly roll up into a pair of coherent, counter-rotating vortices (Fig. 2) that trail behind the aircraft for miles [22].

The evolution of these vortices is a complex, multi-scale fluid dynamic problem governed by two main phases:

- (1) **Descent and Decay:** In free air, the vortex pair descends due to mutual induction and slowly decays due to atmospheric turbulence and fluid instabilities (e.g., Crow instability) [34].
- (2) **Ground Interaction:** During landing, vortices descend into the ground effect zone. This boundary interaction causes the trajectory to diverge hyperbolically [17]. Furthermore, the no-slip condition at the runway surface generates secondary vorticity, which can cause the primary vortices to rebound upward or hover over the runway [19, 36, 49].

This non-linear interaction with the ground, combined with cross-wind variability, makes analytical prediction of vortex trajectories notoriously difficult. Consequently, data-driven approaches that can model these spatio-temporal dynamics from measurement data are highly desirable.

### 3 Related Work

#### 3.1 Wake Vortex Analysis and Forecasting

Early data-driven approaches to wake vortex analysis applied traditional machine learning methods, such as Support Vector Machines (SVMs), random forests, and clustering algorithms to handcrafted features extracted from LiDAR scans [13, 27, 29, 30]. While these methods demonstrated the feasibility of automated wake vortex detection, they relied on superficial scan characteristics and lacked the capacity to represent the complex, multiscale structure of wake vortices. With the advent of deep learning, LiDAR scans were increasingly converted into 2D images and processed using Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM) networks, or hybrid CNN–LSTM architectures for wake vortex detection, segmentation, and localization [1, 2, 4, 8, 43, 47]. More recently, advanced computer vision models such as YOLO and VGG have further improved detection performance on image-based representations of LiDAR data [26, 28, 38, 39].

Despite their success, image-based pipelines require substantial preprocessing and inherently discard the native geometric structure of LiDAR measurements, limiting physical interpretability and explainability [38, 43]. To address this limitation, recent work has begun to operate directly on LiDAR-derived 3D point clouds, which preserve spatial relationships and allow models to exploit the full structure of the measurements. Point cloud-specific architectures such as PointNet [31], PointNet++ [32], and Dynamic Graph Convolutional Neural Networks (DGCNN) [41] have demonstrated strong performance in segmentation and classification tasks across multiple domains [15]. Building on this paradigm, prior work introduced the first application of 3D point cloud semantic segmentation for wake vortex detection, leveraging DGCNN-based models combined with clustering and perturbation-based explanations to achieve accurate and interpretable vortex localization directly from LiDAR scans [33].

However, existing learning-based approaches, both image-based and point cloud-based, remain fundamentally static, processing each LiDAR scan independently and ignoring the strong temporal coherence of wake vortex evolution. As a result, these methods are limited to instantaneous detection or localization and do not explicitly model wake vortex trajectories. This represents a critical gap for operational Wake Vortex Advisory Systems, where decision-making depends not only on identifying current vortices but also on forecasting their future motion and dissipation to safely optimize aircraft separation standards. Moreover, current approaches rely exclusively on supervised learning and require densely annotated scans, which are expensive to obtain and severely limit scalability to large operational datasets.

#### 3.2 Deep Learning on 4D Point Clouds

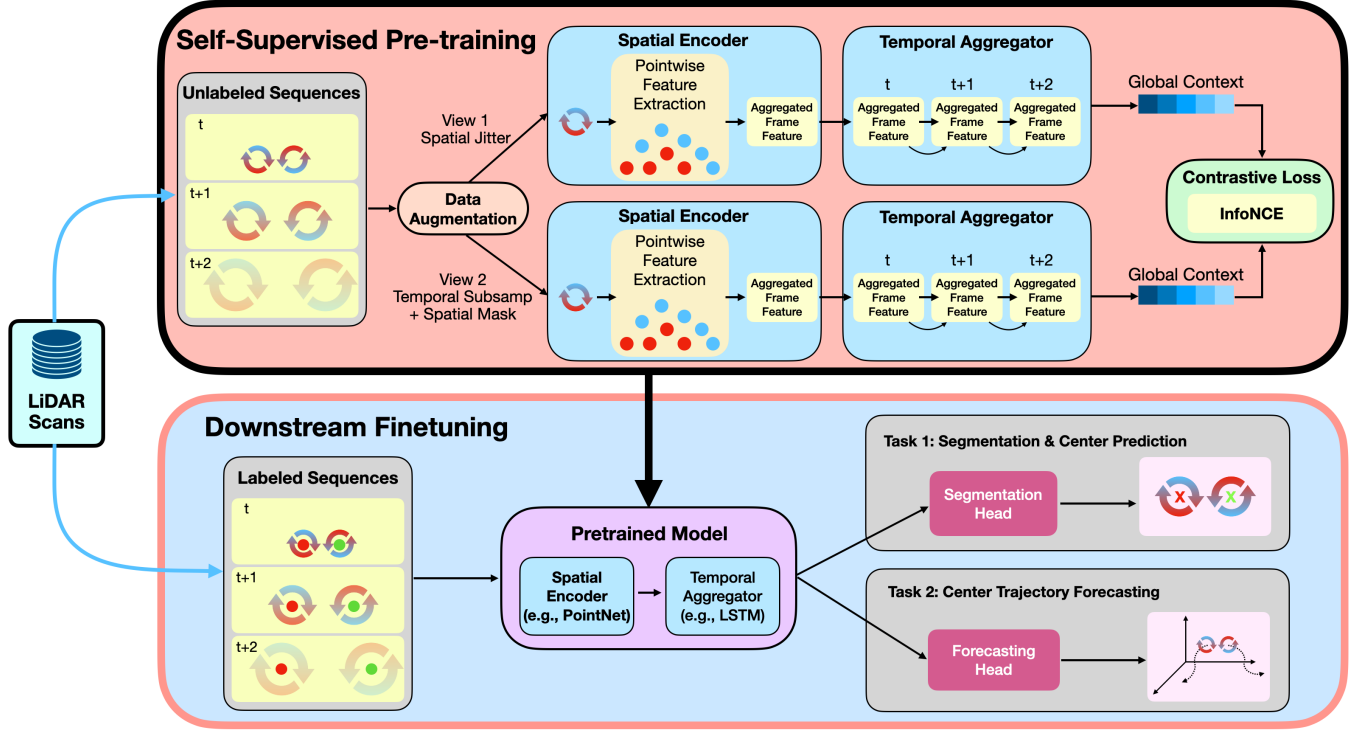
Learning from sequences of point clouds (i.e., 4D data with time as an explicit dimension) has emerged as an active research area, primarily driven by applications in autonomous driving, robotics, and human motion analysis. Early approaches extended static point cloud encoders by aggregating temporal information using Recurrent Neural Networks (RNNs), enabling sequential feature propagation across frames [10, 24]. While effective for short temporal contexts, these methods often struggle to model long-range temporal dependencies and complex spatio-temporal interactions.

More recent architectures explicitly model space and time jointly. Sparse 4D convolutional networks, such as Minkowski-based 4D ConvNets, extend voxel-based convolutions to the spatio-temporal domain and enable efficient learning on large-scale point cloud sequences [7]. Point-based methods have also introduced spatio-temporal convolution operators to directly capture local motion patterns across consecutive frames [12]. In parallel, Transformer-based models leverage self-attention mechanisms to model long-range dependencies in 4D point clouds, achieving strong performance on dynamic scene understanding tasks [11]. These approaches have further been extended to temporally consistent semantic and panoptic segmentation of LiDAR sequences [3].

Despite their expressive power, existing 4D point cloud models are largely designed for scenarios involving rigid or piecewise-rigid motion, such as vehicles, pedestrians, or articulated human bodies, where object identity and geometric structure are preserved over time [3, 7]. In contrast, aircraft wake vortices constitute non-rigid, transient fluid structures whose geometry continuously deforms, diffuses, and dissipates under atmospheric conditions. This fundamental difference limits the direct applicability of existing 4D point cloud methods. Our work addresses this gap by adapting 4D representation learning to wake vortex dynamics, employing a time-distributed architecture that explicitly separates spatial vortex geometry from temporal evolution and decay processes, enabling both detection and trajectory forecasting.

#### 3.3 Self-Supervised Representation Learning

Self-supervised learning (SSL), and contrastive learning in particular, has become a dominant paradigm for representation learning by enabling models to extract semantic structure from large-scale unlabeled data [6, 18]. From a theoretical perspective, contrastive



**Figure 3: Overview of X-VORTEX. Top: self-supervised pre-training on unlabeled wake sequences using weak and strong augmentations of the same sequence, optimized with InfoNCE. Bottom: downstream adaptation on labeled sequences, where the pretrained spatial encoder and temporal aggregator are frozen to produce representations for (i) vortex center localization via a segmentation-based soft-center head and (ii) short-horizon trajectory forecasting via a lightweight prediction head.**

learning has been shown to promote both alignment between positive pairs and uniformity of representations over the embedding space [40]. In the 3D domain, contrastive approaches such as PointContrast [45] and DepthContrast [48] learn invariant representations by enforcing consistency across augmented views of the same underlying geometry, typically at the point or scene level. While effective, most existing 3D SSL methods focus on static scenes, where two views correspond to the same geometry under different spatial transformations.

Extending contrastive learning to temporally evolving data introduces additional challenges, as semantic consistency must be maintained despite genuine structural change over time. Naively contrasting temporally adjacent frames risks suppressing meaningful dynamics, while treating distant frames as negatives may violate underlying continuity. Recent theoretical work on *augmentation overlap* [42, 46] explains how strong but semantically consistent augmentations generate overlapping views of the same sample, enabling contrastive learning to capture shared structure across transformations without collapsing meaningful variation.

Building on these insights, X-VORTEX adopts a spatio-temporal contrastive learning objective tailored to fluid-dynamic point cloud sequences. By treating temporally non-adjacent observations of the same wake vortex and spatially masked regions as positive pairs, the model is encouraged to learn representations that are invariant to decay and deformation while remaining sensitive to physically

meaningful evolution. This enables self-supervised learning of wake vortex dynamics without requiring dense temporal annotations.

## 4 Methodology

### 4.1 Problem Formulation

We formulate the wake vortex analysis problem as learning from a sequence of unstructured 3D point clouds  $\mathcal{S} = \{\mathbf{P}_t\}_{t=1}^T$ , where each frame  $\mathbf{P}_t \in \mathbb{R}^{N \times 3}$  represents a LiDAR scan at time  $t$ , normalized to a fixed number of points  $N$ . The sequence captures the temporal evolution of a vortex pair (port and starboard) generated by an aircraft. Unlike traditional static segmentation, we aim to model the spatio-temporal dynamics of vortex evolution. Our objective is twofold:

- (1) **Unsupervised Representation Learning:** Learn a physics-aware embedding  $Z_S$  from a large corpus of unlabeled sequences by exploiting temporal continuity and spatial invariance.
- (2) **Downstream Detection & Forecasting:** Utilize the learned representation to (a) robustly predict the vortex center coordinates  $(y_c, z_c)$  at time  $t$ , and (b) forecast the future trajectory at  $t + 1, \dots, t + k$ .

## 4.2 The X-VORTEX Framework

We propose **X-VORTEX** (see Fig. 3), a contrastive learning framework grounded in Augmentation Overlap Theory [42, 46]. The theory shows that contrastive learning is effective when augmented views of the same sample retain sufficient distributional overlap while remaining separable from other samples. X-VORTEX leverages this principle to learn robust representations from unlabeled wake vortex sequences. The architecture consists of three core components:

1. *Spatio-Temporal Encoder*. To process 4D data efficiently, we employ a time-distributed architecture:

- **Spatial Backbone (e.g., PointNet)**: Each frame  $P_t$  in the sequence is independently processed by a shared PointNet encoder  $f_\theta$ , yielding a sequence of spatial feature vectors  $\{v_1, \dots, v_T\}$ , where  $v_t \in \mathbb{R}^{D_s}$ .
- **Temporal Aggregator (e.g., LSTM)**: The sequence of spatial features is fed into a Long Short-Term Memory (LSTM) network  $g_\phi$  to capture the vortex decay dynamics. The final hidden state  $\mathbf{h}_T \in \mathbb{R}^{D_t}$  serves as the global context vector for the entire sequence.

2. *Augmentation Strategy (The "Two Views")*. We construct positive pairs by generating two distinct views of the same underlying flight event, specifically designed to enforce physics-based constraints:

- **View 1 (Weak Augmentation)**: The original sequence with minor spatial jitter. This preserves the exact temporal structure.
- **View 2 (Strong Augmentation)**: A perturbed version constructed via:
  - *Temporal Subsampling*: Frames are randomly skipped (e.g.,  $t, t+2, t+4$ ) to force the model to infer missing intermediate states.
  - *Spatial Masking*: Random point dropout (20-30%) and rigid rotation around the vertical axis ( $z$ ). This forces the encoder to rely on global geometry rather than local sensor artifacts.

3. *Contrastive Objective*. We employ the InfoNCE loss to align representations of two augmented views of the same sequence. For each sample, a weakly perturbed anchor view and a strongly augmented counterpart (temporal subsampling and spatial masking) form a positive pair, while views from other sequences in the batch act as negatives. For a batch of  $B$  sequences, the loss for the  $i$ -th sample is:

$$\mathcal{L}_i = -\log \frac{\exp(\text{sim}(\mathbf{h}_i^{(1)}, \mathbf{h}_i^{(2)})/\tau)}{\sum_{j=1}^{2B} \mathbb{1}_{[j \neq i]} \exp(\text{sim}(\mathbf{h}_i^{(1)}, \mathbf{h}_j)/\tau)}$$

where  $\mathbf{h}^{(1)}$  and  $\mathbf{h}^{(2)}$  are the projected embeddings of the two views, and  $\tau$  is the temperature parameter. This objective pulls the strongly perturbed view toward its corresponding anchor, encouraging the model to learn representations that are invariant to spatial sparsity and temporal gaps while preserving the underlying wake dynamics. The weak view retains fine-grained geometric structure, whereas the strong view forces the model to infer wake evolution from

incomplete spatial and temporal observations, mirroring real-world LiDAR sensing conditions.

## 4.3 Downstream Task Adaptation

After pre-training on unlabeled data, we freeze the spatial backbone and fine-tune lightweight heads on the labeled dataset:

*Task 1: Context-Aware Center Prediction*. We introduce a Soft-Center Segmentation Head. Instead of direct regression, we predict a "vortexness" score map  $M \in [0, 1]^N$  for the points in the final frame. The center is computed as the differentially weighted center of mass:

$$\mathbf{c}_{pred} = \frac{\sum_{i=1}^N M_i \cdot \mathbf{p}_i}{\sum_{i=1}^N M_i}$$

This allows the supervision signal (MSE on center coordinates) to backpropagate directly into the segmentation mask, refining the spatial localization.

*Task 2: Trajectory Forecasting*. A Multi-Layer Perceptron (MLP) takes the global context vector  $\mathbf{h}_T$ , computed from a brief history of the three most recent LiDAR scans, and regresses the vortex center coordinates for future time steps  $t+1$  and  $t+2$ .

## 5 Experimental Setup

### 5.1 Dataset and Pre-processing

We utilize a large-scale LiDAR dataset of aircraft wake vortices collected at Vienna International Airport, comprising a total of 1,032,694 raw scans [20].

**Labeled Subset**. The labeled portion consists of 1,176 flight sequences (19,354 scans) with annotated vortex centers, where sequence lengths range from 1 to 52 scans. To ensure consistent temporal context, we retain only sequences with at least five scans; longer sequences are either split into multiple segments or truncated to a fixed length of  $T=5$ . This yields 3,425 sequences containing 17,125 scans in total. The data are split into train/validation/test sets with 2,398/685/342 sequences (11,990/3,425/1,710 scans), respectively.

Ground-truth annotations are provided only for vortex center locations and are generated using a physics-based wake vortex characterization algorithm, the Radial Velocity (RV) method [35]. These labels are subject to inherent uncertainty, particularly under turbulent atmospheric conditions, and recent studies have quantified non-negligible localization errors in such scenarios [44]. We therefore regard these annotations as approximate, algorithm-generated labels rather than exact ground truth.

**Unlabeled Subset (for Pre-training)**. For self-supervised pre-training, we construct an unlabeled dataset of 194,864 sequences (974,320 scans) by grouping consecutive scans with  $\Delta t < 8.0$  s and filtering for valid point densities (at least 8,000 points per scan) to ensure data quality. This subset is split into train/validation/test partitions using a 70%/20%/10% ratio, resulting in 136,406/38,972/19,486 sequences (682,030/194,860/97,430 scans), respectively.

*Sequence Construction*. To strictly prevent data leakage and false negatives in contrastive learning, we employ a Non-Overlapping Chunking strategy. Long flight recordings are split into fixed-length segments of  $T=5$  frames with no overlap. We apply *sequence-level centering* to remove global position and altitude offsets: for each sequence  $\mathcal{S} = \{P_t\}_{t=1}^T$ , we compute a single centroid over all points

across all frames,  $\mu_S = \frac{1}{TN} \sum_{t=1}^T \sum_{i=1}^N \mathbf{p}_{t,i}$ , and translate every point as  $\mathbf{p}'_{t,i} = \mathbf{p}_{t,i} - \mu_S$ . This normalization preserves relative geometry and motion within the sequence while preventing the model from exploiting absolute height or site-specific coordinate offsets.

## 5.2 Baselines

We compare X-VORTEX against representative baselines covering heuristic methods, image-based pipelines, static 3D point cloud models, and temporal models without self-supervised pre-training.

*Heuristic Baselines.* **DBSCAN Clustering** estimates vortex centers as centroids of density-based clusters, while **Intensity Centroid** computes weighted centroids of high-magnitude radial velocity returns. These methods require no learning.

*2D Image-Based Baselines.* LiDAR scans are projected into 2D images and processed using a **Standard CNN** and **YOLOv8**. Predicted vortex regions are converted to center coordinates in physical space.

*Static 3D Point Cloud Baselines.* **PointNet**, **PointNet++**, and **DGCNN** operate directly on 3D point clouds but process each scan independently, without temporal context.

*Temporal Baselines.* A **Supervised LSTM** combines a time-distributed PointNet++ encoder with an LSTM and is trained end-to-end using labeled data only.

*Trajectory Forecasting Baselines.* A **Trajectory-Only LSTM** predicts future vortex centers using past center coordinates, without access to point cloud geometry. We further include two kinematic baselines, **Constant Velocity** and a **Kalman Filter**, which extrapolate future positions based on simple motion assumptions.

## 5.3 Implementation Details

*Architecture.* Unless otherwise stated, X-VORTEX employs PointNet as the spatial encoder and a single-layer LSTM with hidden dimension 256 as the temporal aggregator. The LSTM processes per-frame embeddings and produces a sequence-level representation, which is passed through a two-layer MLP projection head during contrastive pre-training.

*Pre-training protocol.* We pre-train on the unlabeled subset for up to 100 epochs with a global batch size of 64 sequences (128 augmented views per step) using Adam with an initial learning rate of  $10^{-3}$  and a cosine decay schedule. Training is performed on 4 NVIDIA H100 GPUs, with wall-clock time ranging from approximately 40 to 72 hours depending on the backbone architecture. In practice, models typically converge before the full 100 epochs (around 80 epochs). We use the InfoNCE objective with temperature  $\tau=0.07$  and standard in-batch negatives. To avoid false negatives due to scan overlap across training examples, we employ the non-overlapping chunking strategy described in Sec. 5.1.

*Augmentations.* We follow the two-view design in Sec. 4. View 1 applies weak perturbations (minor spatial jitter). View 2 applies strong augmentations: temporal subsampling (random frame skipping within each  $T=5$  chunk), spatial masking via random point dropout with ratio  $p=0.3$ , and rigid rotation around the vertical axis with range  $\pm 30^\circ$ . These parameters are fixed across all experiments for fair comparison.

*Downstream fine-tuning.* For center localization, we freeze the pre-trained spatial encoder and fine-tune the temporal module together with task-specific prediction heads on the labeled subset.

For trajectory forecasting, we train an MLP regressor on top of the sequence-level embedding using the same labeled split.

## 5.4 Evaluation Metrics

*Center Localization Error.* We measure center localization quality using Euclidean RMSE (in meters) between predicted and ground-truth vortex centers in the final frame of each sequence. This metric directly reflects operational localization accuracy.

*Forecasting Error.* For trajectory forecasting, we report RMSE (meters) at horizons  $t+1$  and  $t+2$ , corresponding to approximately 6-8 s and 12-16 s ahead in our sequence sampling.

*Representation Quality and Geometry.* We evaluate unsupervised representation quality via linear probing accuracy on the labeled set. To diagnose contrastive training behavior, we also report alignment and uniformity metrics to quantify how well positive pairs are pulled together while avoiding embedding collapse.

## 6 Results

We evaluate X-VORTEX along three dimensions: (i) representation quality learned from unlabeled sequences, (ii) label efficiency for center localization under scarce supervision, and (iii) the ability to forecast future wake trajectories. We further provide an ablation study to isolate the contributions of key design choices, and qualitative evidence to illustrate typical failure modes and robustness.

**Table 1: Linear Probing Accuracy. Classification accuracy (%) of a linear classifier trained on frozen representations for aircraft type classification (16 classes).**

Method	Backbone	Input View	Acc (%)
Random Initialization	PointNet++	Frame	14.21
Spatial-Only	PointNet++	Frame	29.39
<b>X-VORTEX (Ours)</b>	<b>PointNet</b>	Sequence	<b>72.03</b>
X-VORTEX (Ours)	PointNet++	Sequence	69.41
X-VORTEX (Ours)	DGCNN	Sequence	67.02

### 6.1 Unsupervised Representation Quality

We evaluate representation quality using a standard **linear probing** protocol: the encoder is frozen and a linear classifier is trained on top to predict aircraft type (16 classes) from the labeled subset. While aircraft classification is not a downstream task of interest per se, it provides a diagnostic measure of how linearly separable different wake patterns become under the learned embedding.

As shown in Table 1, X-VORTEX substantially improves linear separability compared to random initialization and spatial-only baselines. In particular, sequence-level representations learned with X-VORTEX achieve up to 72.03% accuracy, whereas single-frame PointNet++ features trained without temporal context reach only 29.39%. This gap indicates that incorporating temporal information during self-supervised pre-training yields embeddings that better represent wake-induced flow structures.

To further analyze representation geometry beyond linear probing, we report alignment and uniformity metrics in Appendix A, providing complementary evidence of stable contrastive learning and well-distributed embeddings.

**Table 2: Center Localization Error. RMSE (meters) of predicted wake vortex centers under different fractions of labeled training data.  $\Delta$  1% denotes the relative improvement of X-VORTEX over the best baseline at 1% labeled data.**

Method	Labeled Data Fraction			$\Delta$ 1%
	1%	10%	100%	
<i>Heuristics (No Learning)</i>				
DBSCAN Clustering	117.45 m	117.45 m	117.45 m	–
Intensity Centroid	118.77 m	118.77 m	118.77 m	–
<i>2D Image Baselines (Projected)</i>				
Standard CNN	56.78 m	35.82 m	30.10 m	–
YOLOv8	27.61 m	25.23 m	22.44 m	–
<i>Static 3D Baselines (Supervised)</i>				
PointNet (Single Frame)	43.75 m	24.54 m	10.18 m	–
PointNet++ (Single Frame)	42.18 m	33.33 m	11.93 m	–
DGCNN (Single Frame)	49.13 m	29.19 m	6.33 m	–
<i>Spatio-Temporal Models (4D)</i>				
Supervised PointNet + LSTM	53.14 m	34.60 m	17.19 m	–
<i>X-VORTEX (finetuned)</i>				
X-VORTEX (PointNet)	<b>9.15 m</b>	<b>7.48 m</b>	<b>5.15 m</b>	<b>+66.8%</b>
X-VORTEX (PointNet++)	15.71 m	10.62 m	7.37 m	+43.1%
X-VORTEX (DGCNN)	20.02 m	17.88 m	10.19 m	+27.5%

## 6.2 Label Efficiency in Center Localization

We evaluate vortex center localization using RMSE (meters) under varying fractions of labeled training data (1%, 10%, and 100%). Results are reported in Table 2. For heuristic methods, which do not depend on labeled training data, the same test-set error is reported across all label fractions. The  $\Delta$  1% column denotes the relative improvement of each X-VORTEX variant over the strongest non-self-supervised baseline (YOLOv8) at 1% labeled data.

Across all settings, X-VORTEX consistently outperforms heuristic, image-based, and fully supervised point-cloud baselines. In the extreme low-label regime (1%), X-VORTEX with a PointNet backbone achieves 9.15 m RMSE, corresponding to a 66.8% reduction in error relative to YOLOv8. Even with alternative backbones, X-VORTEX maintains substantial gains, improving over the best baseline by 43.1% (PointNet++) and 27.5% (DGCNN), demonstrating that the proposed spatio-temporal pre-training strategy generalizes across architectures.

Static 3D models trained from scratch degrade severely under limited supervision, with errors exceeding 40 m at 1% labels, highlighting the difficulty of learning reliable wake structure from sparse annotations alone. In contrast, X-VORTEX leverages large-scale unlabeled sequences to learn transferable representations prior to fine-tuning, enabling accurate localization with minimal labeled data. As label availability increases, X-VORTEX continues to outperform competing methods, reaching 5.15 m RMSE at full supervision.

We observe that the PointNet backbone consistently outperforms PointNet++ and DGCNN across localization and forecasting tasks. This behavior aligns with the physical characteristics of wake vortices, which manifest as diffuse, non-rigid flow structures whose discriminative cues are primarily global (e.g., large-scale velocity dipoles and circulation patterns), rather than localized geometric primitives. PointNet’s global pooling therefore better captures these vortex-scale signatures, while neighborhood-based models rely on local graph constructions that become unstable under LiDAR sparsity, spatial masking, and temporal deformation. This effect is further amplified by our contrastive augmentations, which perturb

local neighborhoods while preserving global structure, making simpler global encoders more robust in the self-supervised setting.

*Discussion.* These results indicate that incorporating temporal context through self-supervised pre-training provides a strong initialization for wake vortex localization. By exploiting unlabeled sequence data rather than relying solely on labeled scans, X-VORTEX substantially reduces annotation requirements while maintaining competitive performance under full supervision.

## 6.3 Wake Vortex Trajectory Forecasting

Beyond localizing current wake positions, operational systems benefit from anticipating where hazardous wakes will move next. We evaluate short-horizon forecasting at  $t+1$  and  $t+2$  using a brief history of 3 past scans, which provides minimal temporal context while reflecting realistic sensing constraints in operational settings.

Results are summarized in Table 3. Kinematic baselines based on Constant Velocity and Kalman filtering exhibit large errors at both horizons, reflecting their inability to capture the non-linear descent and decay of wake vortices. A Trajectory-Only LSTM, which operates solely on past center coordinates, provides only marginal improvement over these methods, indicating that center trajectories alone are insufficient for reliable prediction.

In contrast, X-VORTEX substantially reduces forecasting error across all backbones by leveraging sequence-level representations learned from full point cloud inputs. The PointNet variant achieves the lowest RMSE, with 19.99 m at  $t+1$  and 22.06 m at  $t+2$ , while PointNet++ and DGCNN yield slightly higher but still competitive performance. The growing performance gap at the longer horizon highlights the benefit of incorporating spatial structure when extrapolating wake motion under increasing uncertainty.

*Discussion.* These results indicate that effective wake forecasting requires more than modeling past center trajectories or applying simple kinematic assumptions. By combining temporal context with point-cloud geometry through self-supervised pre-training, X-VORTEX learns representations that support more stable short-term prediction, providing a stronger foundation for forecasting than coordinate-only approaches.

**Table 3: Trajectory Forecasting Error. RMSE (meters) for predicting wake vortex center positions at future time horizons  $t+1$  and  $t+2$ .**

Method	Horizon $t+1$	Horizon $t+2$
Constant Velocity	54.56 m	58.27 m
Kalman Filter	53.83 m	56.57 m
Trajectory-Only LSTM	53.11 m	55.23 m
<b>X-VORTEX (PointNet)</b>	<b>19.99 m</b>	<b>22.06 m</b>
X-VORTEX (PointNet++)	22.31 m	24.76 m
X-VORTEX (DGCNN)	25.37 m	32.79 m

## 6.4 Ablation Studies

We ablate key components of X-VORTEX to assess their contribution to both center localization and trajectory forecasting. Results are reported in Table 4.

*Temporal subsampling is essential.* Removing temporal subsampling substantially degrades performance, increasing center localization error by 13.83 m and forecasting error at  $t+1$  by 9.93 m. This

**Table 4: Ablation Results. RMSE for center localization and trajectory forecasting at horizon  $t+1$  (meters).  $\Delta$  denotes absolute degradation relative to the full X-VORTEX model.**

Configuration	Center ( $\Delta$ )	Horizon $t + 1$ ( $\Delta$ )
X-VORTEX (PointNet)	5.15 m	19.99 m
<i>Augmentations</i>		
w/o Temporal Subsampling (View 2)	18.98 m (+13.83)	29.92 m (+9.93)
w/o Spatial Masking (View 2)	16.99 m (+11.84)	28.84 m (+8.85)
<i>Normalization &amp; Architecture</i>		
w/o Altitude Centering (Raw coords)	10.91 m (+5.76)	25.50 m (+5.51)
w/o LSTM (Mean Pooling)	30.42 m (+25.27)	52.79 m (+32.80)

confirms that forcing the model to bridge temporal gaps during pre-training is critical for learning representations that support both localization and prediction.

*Spatial masking improves robustness.* Disabling spatial masking leads to consistent performance drops (+11.84 m for localization and +8.85 m for forecasting), indicating that invariance to point sparsity and missing returns plays an important role in stabilizing downstream predictions.

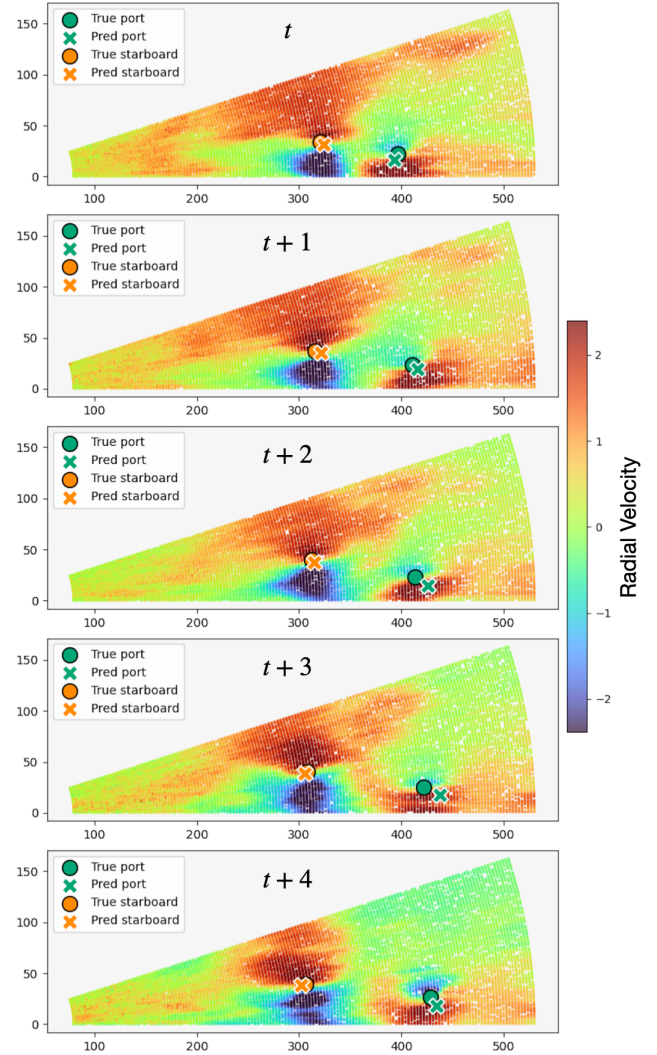
*Sequence-level centering prevents shortcut learning.* Removing altitude centering increases localization and forecasting errors by 5.76 m and 5.51 m, respectively, suggesting that normalization is necessary to prevent reliance on absolute scan placement and to encourage learning of relative wake structure.

*Explicit temporal modeling matters beyond pooling.* Replacing the LSTM with mean pooling causes the largest degradation (+25.27 m for localization and +32.80 m for forecasting), demonstrating that ordered temporal modeling is essential and cannot be replaced by unordered aggregation of frame-level features.

## 6.5 Qualitative Results

Figure 4 shows a representative sequence of five consecutive LiDAR scans ( $t$  to  $t+4$ ) together with ground-truth and predicted vortex centers. Each panel visualizes the radial velocity field measured by the Doppler LiDAR, where warm colors (red/orange) indicate motion toward the sensor and cool colors (blue) indicate motion away from the sensor. The wake consists of a counter-rotating vortex pair, and each vortex therefore exhibits adjacent regions of positive and negative radial velocity, corresponding to the port and starboard vortices trailing the aircraft. Due to their opposite rotation directions, the two vortices display complementary velocity patterns: one typically shows stronger toward-sensor motion in its upper region and away-sensor motion below, while the other exhibits the inverse structure. These characteristic signatures provide important geometric cues for localization but become progressively weaker and more diffuse as the wake decays over time. Green and orange circles denote ground-truth centers for the port and starboard vortices, respectively, while crosses indicate model predictions.

Despite this progressive decay and increasing sparsity, X-VORTEX maintains stable center predictions across the sequence, closely tracking both vortices even when wake signatures are faint. This highlights the benefit of jointly leveraging temporal context and spatial structure: rather than relying on isolated high-contrast regions in individual scans, X-VORTEX integrates information across time to preserve coherent localization.



**Figure 4: Qualitative Visualization.** A sequence of five consecutive LiDAR scans ( $t$  to  $t+4$ ) showing the radial velocity field, where warm colors indicate motion toward the sensor and cool colors indicate motion away. The counter-rotating port and starboard vortices appear as paired regions of opposite velocity. Green and orange circles denote ground-truth centers, while crosses indicate model predictions.

Qualitative performance is especially relevant in operational settings, where the most challenging cases arise not during early, high-contrast wake stages, but during later phases when vortices are partially dissipated yet still potentially hazardous. The consistency of X-VORTEX in these borderline conditions aligns with its design goal of learning representations robust to wake decay and sensor sparsity. Additional qualitative examples across diverse wake conditions are provided in Appendix B.

## 7 Conclusion

We introduced X-VORTEX, a self-supervised spatio-temporal contrastive learning framework for wake vortex analysis from LiDAR

point cloud sequences. By leveraging *Augmentation Overlap Theory* and treating temporal evolution as a natural augmentation, X-VORTEX learns transferable sequence-level representations from large-scale unlabeled data. Experiments on nearly one million scans demonstrate that X-VORTEX substantially reduces labeling requirements for vortex center localization and enables accurate short-horizon trajectory forecasting, consistently outperforming heuristic, image-based, and fully supervised baselines. Ablation studies further confirm the importance of temporal subsampling, spatial masking, sequence-level normalization, and explicit temporal modeling in achieving these gains.

Future work will explore extending X-VORTEX to longer-horizon forecasting with uncertainty estimation, integrating continuous-time temporal models, and incorporating meteorological context such as wind and turbulence profiles. Beyond wake vortices, we believe the proposed spatio-temporal self-supervised framework can be applied to other atmospheric flow phenomena captured by remote sensing, opening new opportunities for data-driven airspace safety and efficiency.

## Acknowledgments


The authors acknowledge the financial support by the Federal Ministry of Research, Technology, and Space of Germany and by the Saxon State Ministry for Science, Culture, and Tourism, in the programme Center of Excellence for AI-research „Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig“, project identification number: ScaDS.AI.

The authors acknowledge the computing time made available to them on the high-performance computer at the NHR Center of TU Dresden. This center is jointly supported by the Federal Ministry of Research, Technology and Space of Germany and the state governments participating in the NHR ([www.nhr-verein.de/unsere-partner](http://www.nhr-verein.de/unsere-partner)).

## References

- Jennifer Abras and Nathan S Hariharan. 2021. Application of machine learning to automate flow-physics identification in computed solutions: Hover rotor wake vortex identification and breakdown analysis. In *AIAA Scitech 2021 Forum*. 0474.
- Yi Ai, Yuanji Wang, Weijun Pan, and Dingjie Wu. 2021. A deep learning framework based on multisensor fusion information to identify the airplane wake vortex. *Journal of Sensors* 2021 (2021), 1–10.
- Mehmet Aygun, Aljosa Osep, Mark Weber, Maxim Maximov, Cyrill Stachniss, Jens Behley, and Laura Leal-Taixé. 2021. 4d panoptic lidar segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 5527–5537.
- Nikolay Baranov and Boris Resnick. 2021. Wake vortex detection by convolutional neural networks. *European Journal of Electrical Engineering and Computer Science (EEACS)* 3 (2021), 92–97.
- Bundesstelle für Flugunfalluntersuchung (BFU). 2017. *Interim Report: Serious Incident, Bombardier Challenger 604 vs Airbus A380-800, Arabian Sea*. Technical Report BFU 17-0024-2X. German Federal Bureau of Aircraft Accident Investigation, Braunschweig, Germany.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*. PmlR, 1597–1607.
- Christopher Choy, JunYoung Gwak, and Silvio Savarese. 2019. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 3075–3084.
- Nana Chu, Kam KH Ng, Ye Liu, Kai Kwong Hon, Pak Wai Chan, Jianbing Li, and Xiaoge Zhang. 2024. Assessment of approach separation with probabilistic aircraft wake vortex recognition via deep learning. *Transportation Research Part E: Logistics and Transportation Review* 181 (2024), 103387.
- EUROCONTROL. 2018. European Aviation in 2040. Challenges of Growth. <https://www.eurocontrol.int/sites/default/files/content/documents/official-documents/reports/challenges-of-growth-2018.pdf>
- Hehe Fan and Yi Yang. 2019. PointRNN: Point recurrent neural network for moving point cloud processing. *arXiv preprint arXiv:1910.08287* (2019).
- Hehe Fan, Yi Yang, and Mohan Kankanhalli. 2021. Point 4d transformer networks for spatio-temporal modeling in point cloud videos. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 14204–14213.
- Hehe Fan, Xin Yu, Yuhang Ding, Yi Yang, and Mohan Kankanhalli. 2022. Pstnet: Point spatio-temporal convolution on point cloud sequences. *arXiv preprint arXiv:2205.13713* (2022).
- Melvin A Felton, Prabhat Kumar, Jana Catuche, Deryck James, and David A Ligon. 2025. Characterizing wake vortex pairs using unsupervised machine learning. *Optics Express* 33, 18 (2025), 37795–37813.
- Michael Frech, Frank Holzäpfel, Thomas Gerz, and Jens Konopka. 2002. Short-term prediction of the horizontal wind vector within a wake vortex warning system. *Meteorological Applications* 9, 1 (2002), 9–20.
- Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennamoun. 2021. Deep Learning for 3D Point Clouds: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.* 43, 12 (2021), 4338–4364.
- James N Hallock and Frank Holzäpfel. 2018. A review of recent wake vortex research for increasing airport capacity. *Progress in Aerospace Sciences* 98 (2018), 27–36.
- JK Harvey and Fi J Perry. 1971. Flowfield produced by trailing vortices in the vicinity of the ground. *AIAA journal* 9, 8 (1971), 1659–1660.
- Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. 2020. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 9729–9738.
- Frank Holzäpfel and Meiko Steen. 2007. Aircraft wake-vortex evolution in ground proximity: analysis and parameterization. *American Institute of Aeronautics and Astronautics journal* 45, 1 (2007), 218–227.
- Frank Holzäpfel, Anton Stephan, Grigory Rotshteyn, Stephan Körner, Norman Wildmann, Lothar Oswald, Thomas Gerz, Günther Borek, Alexander Floh, Christian Kern, et al. 2021. Mitigating wake turbulence risk during final approach via plate lines. *AIAA Journal* 59, 11 (2021), 4626–4641.
- Stephan Körner and Frank Holzäpfel. 2019. Assessment of the wake-vortex proximity to landing aircraft exploiting field measurements. *Journal of Aircraft* 56, 3 (2019), 1250–1258.
- Pijush K. Kundu and Ira M. Cohen. 2002. *Fluid Mechanics* (2nd ed.). Academic Press. 629 – 660 pages.
- Jianbing Li, Hang Gao, Tao Wang, and Xuesong Wang. 2017. A survey of the scattering characteristics and detection of aircraft wake vortices. *Journal of Radars* 6, 6 (2017), 660–672.
- Xingyu Liu, Mengyuan Yan, and Jeannette Bohg. 2019. Meteornet: Deep learning on dynamic 3d point cloud sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 9246–9255.
- National Transportation Safety Board. 2004. *In-Flight Separation of Vertical Stabilizer, American Airlines Flight 587, Airbus Industrie A300-605R, N14053, Belle Harbor, New York, November 12, 2001*. Technical Report NTSB/AAR-04/04. National Transportation Safety Board, Washington, D.C.
- Weijun Pan, Yingjie Duan, Qiang Zhang, Jiahao Tang, and Jun Zhou. 2019. Deep learning for aircraft wake vortex identification. *IOP Conference Series: Materials Science and Engineering* 685, 1 (2019), 012015.
- Weijun Pan, Yuanfei Leng, and Shuai Han. 2021. Aircraft wake detection method based on SVM and morphology. In *2021 IEEE 3rd International Conference on Civil Aviation Safety and Information Technology (ICCASIT)*. IEEE, 356–360.
- Weijun Pan, Yuanfei Leng, Haoran Yin, Xiaolei Zhang, et al. 2022. Identification of aircraft wake vortex based on VGGNet. *Wireless Communications and Mobile Computing* 2022 (2022).
- Weijun Pan, Zhengyuan Wu, and Xiaolei Zhang. 2020. Identification of aircraft wake vortex based on SVM. *Mathematical Problems in Engineering* 2020 (2020), 1–8.
- Weijun Pan, Haoran Yin, Yuanfei Leng, and Xiaolei Zhang. 2022. Recognition of aircraft wake vortex based on random forest. *IEEE Access* 10 (2022), 8916–8923.
- Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 652–660.
- Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* 30 (2017).
- Zhan Qu, Shuzhou Yuan, Michael Färber, Marius Brennfleck, Niklas Wartha, and Anton Stephan. 2025. Explainable LiDAR 3D Point Cloud Segmentation and Clustering for Detecting Airplane-Generated Wind Turbulence. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*. 2504–2513.
- Robert E Robins, Donald P Delisi, and George C Greene. 2001. Algorithm for prediction of trailing vortex evolution. *Journal of aircraft* 38, 5 (2001), 911–917.
- Igor N Smalikhov, VA Banakh, Frank Holzäpfel, and Stephan Rahm. 2015. Method of radial velocities for the estimation of aircraft wake vortex parameters from data measured by coherent Doppler lidar. *Optics Express* 23, 19 (2015), A1194–A1207.

- [36] PR Spalart, M Kh Strelets, AK Travin, and ML Shur. 2001. Modeling the interaction of a vortex pair with the ground. *Fluid Dynamics* 36, 6 (2001), 899–908.
- [37] Anton Stephan, David Rohlmann, Frank Holzäpfel, and Ralf Rudnik. 2019. Effects of detailed aircraft geometry on wake vortex dynamics during landing. *Journal of Aircraft* 56, 3 (2019), 974–989.
- [38] Anton Stephan, Grigory Rotshteyn, Niklas Wartha, Frank N Holzäpfel, Nicolas Petross, and Lars Stietz. 2023. Artificial Neural Networks for Individual Tracking and Characterization of Wake Vortices in LIDAR Measurements. In *AIAA AVIATION 2023 Forum*. 3682.
- [39] Lars Olaf Stietz. 2022. *Artificial Neural Networks for Individual Tracking and Characterization of Wake Vortices in LiDAR Measurements*. Ph.D. Dissertation. Universität Hamburg.
- [40] Tongzhou Wang and Phillip Isola. 2020. Understanding contrastive representation learning through alignment and uniformity on the hypersphere. In *International conference on machine learning*. PMLR, 9929–9939.
- [41] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. 2019. Dynamic Graph CNN for Learning on Point Clouds. *ACM Trans. Graph.* 38, 5 (2019), 146:1–146:12.
- [42] Yifei Wang, Qi Zhang, Yisen Wang, Jiansheng Yang, and Zhouchen Lin. 2022. Chaos is a Ladder: A New Theoretical Understanding of Contrastive Learning via Augmentation Overlap. In *International Conference on Learning Representations*.
- [43] Niklas Wartha, Anton Stephan, Frank Holzäpfel, and Grigory Rotshteyn. 2022. Characterizing aircraft wake vortex position and strength using LiDAR measurements processed with artificial neural networks. *Optics Express* 30, 8 (2022), 13197–13225.
- [44] Niklas Wartha, Anton Stephan, Frank N Holzäpfel, and Grigory Rotshteyn. 2023. Investigating Errors of Wake Vortex Retrievals Using High Fidelity Lidar Simulations. In *AIAA AVIATION 2023 Forum*. 3679.
- [45] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. 2020. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European conference on computer vision*. Springer, 574–591.
- [46] Qi Zhang, Yifei Wang, and Yisen Wang. 2025. An Augmentation Overlap Theory of Contrastive Learning. *Journal of Machine Learning Research* 26, 228 (2025), 1–42.
- [47] Xinyu Zhang, Hongwei Zhang, Qichao Wang, Xiaoying Liu, Shouxin Liu, Rongchuan Zhang, Rongzhong Li, and Songhua Wu. 2024. Locating and Grading of Lidar-Observed Aircraft Wake Vortex Based on Convolutional Neural Networks. *Remote Sensing* 16, 8 (2024), 1463.
- [48] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. 2021. Self-supervised pretraining of 3d features on any point-cloud. In *Proceedings of the IEEE/CVF international conference on computer vision*. 10252–10263.
- [49] ZC Zheng and Robert L Ash. 1996. Study of aircraft wake vortex behavior near the ground. *AIAA journal* 34, 3 (1996), 580–589.



**Figure 5: Alignment and uniformity during self-supervised pre-training (PointNet backbone). Validation alignment and uniformity over the first 20 epochs are shown for clarity. Alignment improves rapidly as positive pairs are pulled together, while uniformity decreases steadily, indicating increasing dispersion of embeddings across the hypersphere. Both metrics stabilize early, suggesting stable contrastive training without representation collapse.**

## A Alignment and Uniformity

To better study the behavior of X-VORTEX during self-supervised pre-training, we analyze the learned embedding geometry using the **alignment** and **uniformity** metrics introduced by Wang and Isola [40]. Alignment measures the expected distance between positive pairs, reflecting how tightly augmented views of the same sequence are pulled together. Uniformity measures how evenly embeddings are distributed over the unit hypersphere, penalizing representation collapse and promoting separation between different samples.

Figure 5 reports validation alignment and uniformity over the first 20 pre-training epochs for the PointNet backbone (shown for clarity). Alignment rapidly improves within the first few epochs and then stabilizes, indicating that temporally and spatially augmented views of the same wake sequence are consistently mapped to nearby embeddings. At the same time, uniformity decreases steadily (becoming more negative), reflecting improved dispersion of representations across the embedding space.

Concretely, alignment increases from near zero at initialization to stable values around 0.05, while uniformity drops from approximately 0 to about  $-2.38$ , suggesting that X-VORTEX simultaneously achieves strong intra-sequence invariance and maintains inter-sequence diversity. Importantly, both metrics stabilize early in training, indicating well-behaved contrastive optimization without embedding collapse.

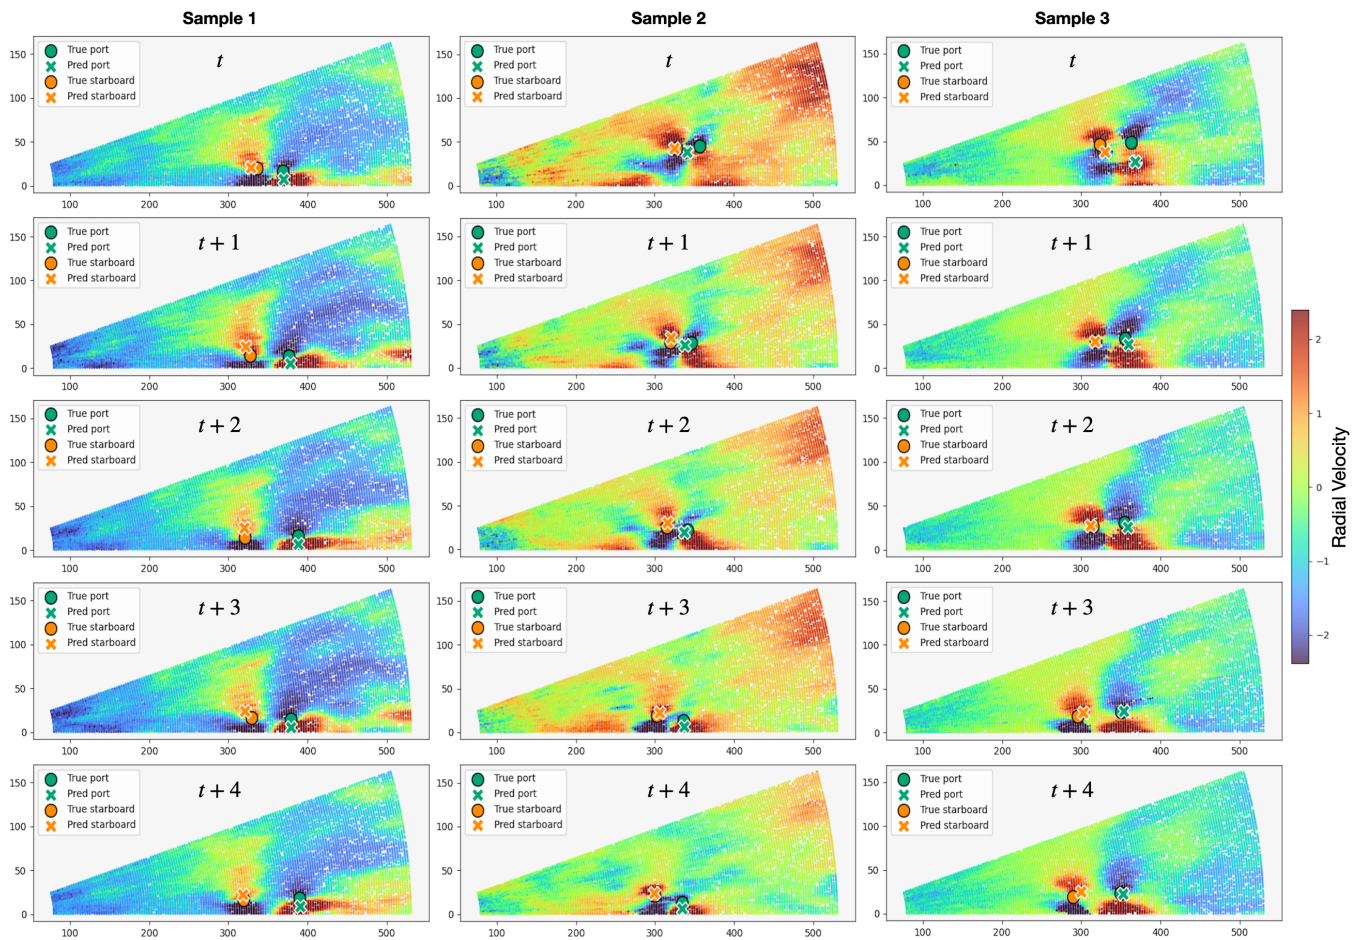
This balance between alignment and uniformity helps explain the strong downstream performance observed with limited labeled data: tight alignment encourages robustness to temporal subsampling and spatial sparsity, while good uniformity preserves discriminability across different wake events. Together, these properties support effective transfer of the learned representations to localization and forecasting tasks.

## B Additional Qualitative Visualizations

This appendix provides additional qualitative examples for both wake vortex center localization and trajectory forecasting, complementing the representative cases shown in Sec. 6.5. Each visualization displays Doppler LiDAR radial velocity fields, where warm colors indicate motion toward the sensor and cool colors indicate motion away. Green and orange circles denote ground-truth centers of the port and starboard vortices, respectively, while crosses indicate model predictions.

*Center Localization.* Figures 6 and 7 show two sets of three test sequences, each consisting of five consecutive scans ( $t$  to  $t+4$ ). These examples cover diverse wake strengths, decay patterns, and background conditions. Despite progressive weakening of velocity contrast and increasing sparsity in later frames, X-VORTEX consistently tracks both vortex cores and maintains stable predictions across time.

*Trajectory Forecasting.* Figure 8 presents three representative forecasting examples. For each sequence, ground-truth centers are shown up to  $t$ , followed by model predictions at  $t+1$  and  $t+2$ . These cases illustrate X-VORTEX’s ability to extrapolate wake motion beyond the observed history, even under non-linear descent and asymmetric decay between the port and starboard vortices.



**Figure 6: Additional center localization examples (Samples 1–3). Each column shows a wake sequence over five consecutive scans ( $t$  to  $t+4$ ). Circles indicate ground-truth vortex centers and crosses denote X-VORTEX predictions.**

Across all samples, X-VORTEX demonstrates robustness to wake dissipation, point sparsity, and background clutter, supporting the quantitative improvements reported in Secs. 6.2 and 6.3.

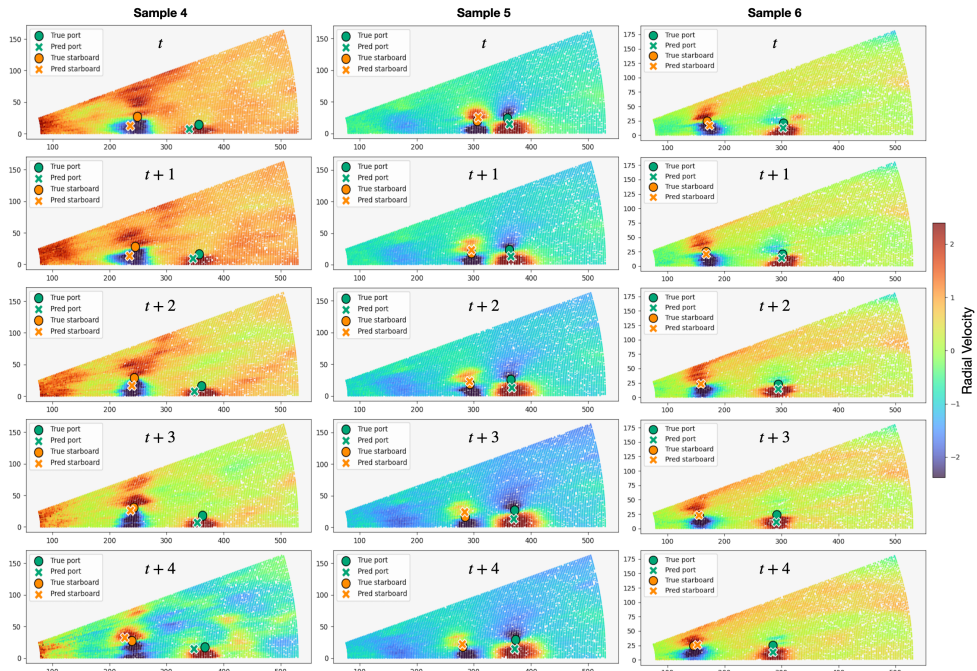


Figure 7: Additional center localization examples (Samples 4–6). X-VORTEX maintains consistent localization across diverse wake conditions and progressive decay.

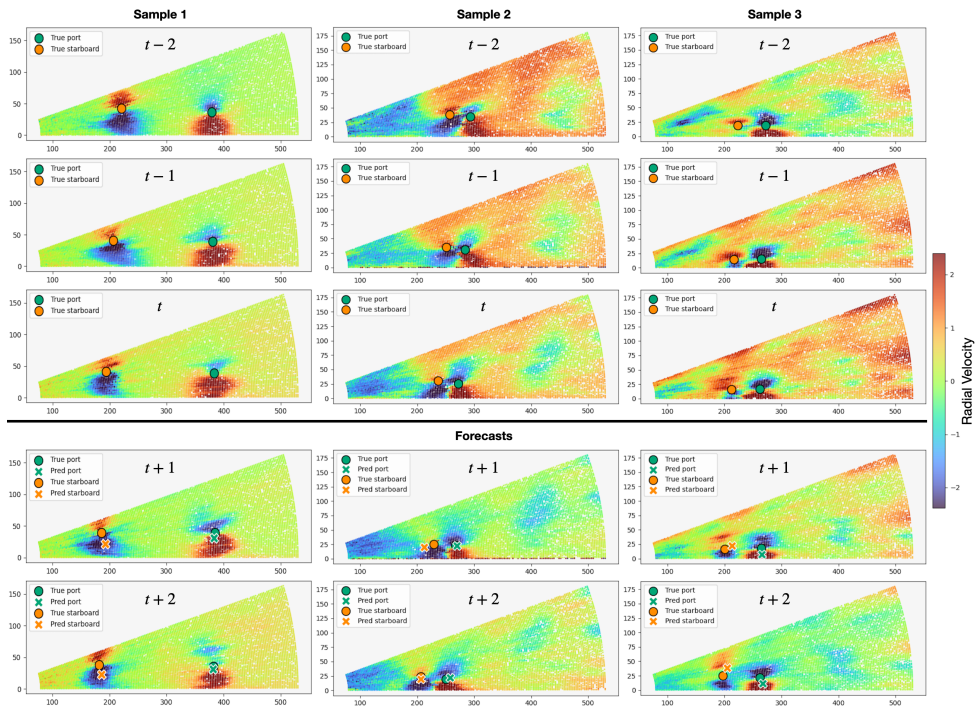


Figure 8: Additional trajectory forecasting examples. Ground-truth centers are shown up to  $t$ , followed by X-VORTEX predictions at  $t+1$  and  $t+2$ , illustrating short-horizon extrapolation under non-linear wake evolution.