

UniField: Joint Multi-Domain Training for Universal Surface Pressure Modeling

Junhong Zou^{1,2} Zhenxu Sun³ Wei Qiu³ Zhaoxiang Zhang^{1,2}
Zhen Lei^{1,2*} Xiangyu Zhu^{1,2*}

¹MAIS, Institute of Automation, Chinese Academy of Sciences, Beijing, China,

²School of Artificial Intelligence, University of Chinese Academy of Sciences, Beijing, China,

³Key Laboratory for Mechanics in Fluid Solid Coupling Systems, Institute of Mechanics, Chinese Academy of Sciences, Beijing, China,

<https://github.com/zoujunhong/UniField>

Abstract

Aerodynamic simulation of the surface pressure field around objects is crucial for many engineering problems. In recent years, deep neural networks have emerged as an efficient alternative to traditional, computationally expensive CFD simulations for modeling surface pressure fields. However, data scarcity remains a fundamental challenge, limiting the application of neural networks. To address this limitation, we propose to integrate aerodynamic data from multiple subfields and conduct joint training to learn more general field representations. We consolidate five different datasets covering various fields, including automobiles, trains, aircraft, and general shapes. Facing significant data differences across different domains, we propose UniField, which employs a domain-agnostic Transformer module to extract general point cloud features and customizes domain-specific flow-conditioned adapters to adapt to the flow information in different subfields. Despite the fact that aerodynamic data from different subfields are typically governed by different equations, we compare models trained jointly on all data with those trained separately on individual datasets and find that the jointly-trained model commonly demonstrates better performance. This indicates that these data complement each other to help the model learn better flow field representations. These results highlight the potential of UniField as a universal flow field representation model and lay the foundation for broader applications of neural networks in aerodynamic analysis.




Figure 1. **Overview.** UniField is a surface pressure foundation model pre-trained on multiple aerodynamic datasets covering diverse geometry and flow regimes (e.g., automobiles, trains, aircraft, and other generic shapes). By jointly learning from these subfields, the model captures unified representation of geometry and flow field, exhibiting improved generalization and achieves better surface pressure predictions across unseen geometries under data scarcity situations.

1. Introduction

Analysis of surface pressure field around objects is crucial in many engineering problems. For instance, in transportation systems such as automobiles, trains, and aircraft, surface pressure fields act as key physical quantities for computing lift, and drag, which is essential for optimizing aerodynamic performance [11], ensuring structural integrity, and enhancing fuel efficiency [6]. Traditionally, these parameters have been simulated using Computational Fluid

*Corresponding authors

Dynamics (CFD), a method that is expensive and time-consuming, making it difficult to scale up to analysis of vast shapes and flow condition combinations.

Deep neural network methods [1, 2, 5, 17] attempt to learn complex flow field patterns from existing CFD data. Compared with CFD simulation, which typically requires hundreds of CPU hours to simulate pressure fields, neural networks can usually complete predictions within seconds, thus becoming a promising and efficient alternative. However, data scarcity remains a fundamental obstacle to the broad deployment of these models in aerodynamic applications. In many specialized subfields, such as trains or aircraft, there are usually only a small amount of available CFD data [11, 21–23], making it difficult to train models with strong generalization capabilities. This limitation greatly restricts the effectiveness of neural network methods. A potential solution is to combine flow field data from different subfields to make up for the insufficient data in a single subfield and train domain-general flow field representation models. However, in different subfields, flow field distributions are usually controlled by different equations, which makes it difficult for models to learn generalizable flow field knowledge. This makes researchers prefer to build domain-specific or even object-specific neural network models, instead of exploring general flow field analysis models.

In this paper, we verify the feasibility of training a universal model through datasets from multiple subfields. First, we collected five datasets from multiple different fields, including an automotive datasets DrivAerNet++ [5], a train dataset [11], an aircraft wing dataset, an full aircraft dataset [17], and the general shape dataset FlowBench [14]. They involve completely different scenarios and flow field conditions: the automotive and train datasets describe the flow field distribution around the vehicle during horizontal movement; the aircraft datasets involve much higher-speed movement states and introduce the angle of attack to describe vertical movement; the general shape dataset FlowBench contains lid-driven cavity flow, which is significantly different from the flow around moving objects described by other datasets. To be able to handle data from multiple subfields simultaneously, we proposed the Unified Field Learning framework (UniField), which can effectively extract common features from the data and customize independent flow field information extraction modules for the specific flow field conditions of each field, thereby achieving joint training on multiple datasets. Specifically, we first represent the surface of the object using a unified point cloud format data and use the point transformer [24] as a universal point cloud feature extraction module to extract domain-general geometric features. On the other hand, to address the significant differences in flow field conditions across different fields, we customize a set of parallel flow-

conditioned adapter, each one for adapting to a subfield. For the input data, we route it to its corresponding flow-conditioned adapter for flow field information extraction. Based on this design, UniField can effectively handle flow field data from different fields and learn universal flow field knowledge.

After joint training on multiple datasets, UniField demonstrates the ability to learn universal flow field knowledge and enhance the performance of each subfield. Specifically, we compared models trained on a single dataset only with those trained on multiple datasets jointly. Under the same settings, the joint-trained models showed performance advantages in all subfields, and this advantage was more pronounced in data-scarce domains. For instance, in the train dataset, which contains only three different train shapes, with two used for training and one for testing, the jointly trained model reduced the error by more than 50% in the test of the train dataset compared to the model trained only on train data. Additionally, the absolute performance of UniField has already reached the current SOTA performance according to the comparison result on the public automotive dataset DrivAerNet++.

To summarize, we propose UniField, a universal framework for analyzing fluid field around objects. UniField achieves cross-domain and cross-dataset joint fluid field model training by constructing a domain-agnostic point cloud feature extraction module and customizing domain-specific flow field information extraction modules for the vastly different flow field conditions in different domains. We compared models trained only on a single dataset with those trained on multiple datasets jointly, and verified that the models trained jointly on all datasets were either better than or comparable to the single-dataset models. Among them, for fields with scarce data, such as the aircraft and train datasets, joint training could significantly reduce the model’s error in the target field. This serves as strong evidence that integrating multiple domain datasets to augment data can effectively alleviate the shortage of CFD data.

2. Related Work

2.1. Surface Pressure Prediction

Surface pressure prediction aims to infer the pressure distribution across a vehicle’s surface given its geometry and flow conditions. Traditional CFD software simulates surface fields by numerically solving the governing equations of fluid dynamics, such as the Navier–Stokes equations, over a discretized computational domain, but is computationally heavy, motivating neural networks as alternatives. DrivAerNet++ [5] is a large-scale aerodynamic dataset. It provides thousands of car geometries, CFD flow and pressure fields, parametric models, and aerodynamic coefficients, fostering robust training and evaluation for automotive aerodynam-




Figure 2. **UniField Architecture.** UniField learns surface pressure fields across multiple subfields, such as automobiles, aircraft, and general geometric shapes. Given point cloud as input, UniField uses a UNet-style network to predict dense pressure fields. Each network layer first extracts the geometric features using a Point Transformer layer, followed by a flow condition adapter (FCA) module that injects domain-specific information through domain routing modulation, thereby achieving cross-domain adaptation. The semantic aggregation module iteratively groups semantically related points through attention and recurrent updates to achieve downsampling.

ics. It also introduces RegDCGNN [4], a dynamic graph convolutional neural network to regress aerodynamic parameters, while avoiding the overhead of rendering or SDF preprocessing. In the subsequent work, Transolver [10, 18] leverages a Transformer-based PDE solver with a physics-inspired slice attention that groups mesh points into learnable physical-state slices, enabling scalable generalization across complex geometries. Factorized Implicit Global Convolution (FIGConvNet) [2] efficiently learns global interactions across 3D meshes via implicit factorization, reducing complexity while preserving accuracy. TripNet [1] encodes 3D car geometry into compact triplane representations, enabling point-wise predictions of pressure and full flow fields.

2.2. Generalization of Neural Models for Flow Field Modeling

Existing deep learning approaches for fluid modeling exhibit substantial differences in their applicability and generalization ability. We summarize these models into the following categories.

- **Models for Single-Scene or Fixed Geometry.** Some models are only designed for Single-Scene or Fixed Geometry. For instance, compressive sensing methods [11] represent specific field data as a linear combination of POD bases, PROSNet [26] optimizes the reconstruction of the field and the arrangement of sensors for a specific high-speed train geometry, and PINNs [12, 19] explicitly embeds partial differential equation (PDE) constraints into the loss function to enforce physical con-

sistency. While effective for a single geometry or under fixed boundary conditions, their model parameters are tightly coupled with specific physical settings, resulting in limited generalization across domains.

- **Neural Operators.** Neural operator approaches [8, 15, 16] aim to learn the mapping between input and output fields [7], enabling fast approximations of specific PDEs. Compared with PINNs, they offer broader applicability, but remain constrained to fixed governing equations, making them difficult to transfer across fundamentally different physical regimes (e.g., incompressible vs. compressible flows, subsonic vs. transonic regimes).
- **Symbolic or Parameterized PDE-Driven Solvers.** Recent frameworks such as UniSolver [25] and PDEformer [20] attempt to introduce the symbolic or parametric representation of PDEs as model inputs, moving toward general PDE solvers. Such methods still rely on complete prior knowledge of the governing equations and precise boundary conditions.
- **Data-Driven Field-Centric Modeling.** A growing line of research [1, 2, 17] seeks to eliminate explicit dependence on predefined PDE forms by reformulating flow-field prediction as a geometry-conditioned regression problem. These models directly take geometric point clouds as inputs to predict target fields such as pressure and velocity. This formulation opens the possibility for building generalizable field foundation models, yet research in this direction remains limited. Our proposed UniField follows this paradigm, leveraging multi-domain data and joint training to achieve cross-geometry and

cross-velocity generalization.

3. Method

This section presents UniField, a multi-domain point-cloud network for surface pressure prediction across heterogeneous aerodynamic domains (e.g., ground vehicles, aircraft, and generic bodies). A **single** geometric backbone is shared by all domains, while a set of **parallel** Flow-Conditioned Adapters specialize to domain-specific flow conditions.

Problem setup. An input sample comprises a surface point cloud $G \in \mathbb{R}^{N \times 3}$ and a source identifier $s \in \{1, \dots, M\}$ specifying one of M domains. Also, each domain s is accompanied by its own flow-condition vector $C^{(s)} \in \mathbb{R}^{D_f^{(s)}}$ (e.g., velocity v and wind speed w for train; Mach number Ma and Angle of Attack AoA for aircraft; and Reynolds number Re for general shape dataset). As the model output, the model predicts a scalar pressure coefficient for every surface node, $P \in \mathbb{R}^N$.

Architecture overview. The overview of UniField is shown in Fig. 2. The network follows an UNet-Style topology tailored to dense point-wise prediction. The encoder stacks vector self-attention blocks to extract geometric features; after each block, features are modulated by a domain-specific flow-conditioned adapters. Downsampling is performed by a semantic aggregation module that groups points by features’ semantic similarity. The decoder restores resolution via kNN interpolation, merges skip connections, and a final linear head outputs point-wise pressure coefficient prediction.

3.1. Vector Self-Attention on Point Clouds

Vector self-attention is proposed in Point Transformer [24] for point cloud processing. Its pattern of calculating attention within the k -nearest neighbor range is suitable for processing point clouds with a large number of points. Let input $= (x, p)$ denote point features $x \in \mathbb{R}^{N \times D}$ and coordinates $p \in \mathbb{R}^{N \times 3}$. Each self-attention layer computes, for the i -th point,

$$\begin{aligned} \delta_{ij} &= \mathcal{P}(p_i - p_j), \\ q_i, k_i, v_i &= \mathcal{Q}(x_i), \mathcal{K}(x_i), \mathcal{V}(x_i), \\ y_i &= \sum_{x_j \in N_k(x_i)} \text{Softmax}(\gamma(q_i - k_j + \delta_{ij})) \odot (v_j + \delta_{ij}), \end{aligned} \quad (1)$$

where $\mathcal{Q}, \mathcal{K}, \mathcal{V}$ are linear projections, and $\mathcal{P}, \gamma$ are two-layer MLPs; $N_k(\cdot)$ denotes the k -nearest neighbors. Subsequently, the features are integrated with the attention output with a residual connection, and finally processed by a feed-forward network (FFN) with two linear layers and a GELU

activation:

$$x_i \leftarrow x_i + y_i, \quad x_i \leftarrow x_i + \text{FFN}(x_i). \quad (2)$$

All self-attention layers and the decoder are shared across all subfields.

3.2. Semantic Aggregation Module

Conventional point cloud downsampling methods (e.g., FPS + kNN pooling) relies on spatial proximity only. To better align with cross-domain geometry, we employ a semantic aggregation module adapted from Slot Attention [9], which clusters points using learnable feature similarity.

Given input $= (x, p)$, we first choose K initial centers $S = (x_s, p_s)$ via FPS, then refine x_s by attention-based aggregation:

$$\begin{aligned} \delta_{ij} &= \mathcal{P}_s(p_{s,i} - p_j), \\ q_i, k_j, v_j &= \mathcal{Q}_s(x_{s,i}), \mathcal{K}_s(x_j), \mathcal{V}_s(x_j), \\ y_{s,i} &= \sum_{x_j \in N_k(x_{s,i})} \text{Softmax}(\gamma_s(q_i - k_j + \delta_{ij})) \odot (v_j + \delta_{ij}), \\ x_s &\leftarrow \text{GRU}(x_s, y_s), \\ x_s &\leftarrow x_s + \text{FFN}(x_s). \end{aligned} \quad (3)$$

Here $\mathcal{Q}_s, \mathcal{K}_s, \mathcal{V}_s$ are linear layer, $\mathcal{P}_s$ and γ_s are two-layer-MLPs, and $N_k(x_{s,i})$ denotes the k -nearest neighbors of slot $x_{s,i}$ within the point set x . GRU represents a gate recurrent unit, while FFN is a feed-forward network.

3.3. Parallel Flow-Conditioned Adapters

Different subfields commonly require distinct flow descriptors (e.g., (v, w) for train vs. (Ma, AoA) for aircraft), which we encode using a bank of parallel flow-conditioned adapters $\{\text{FCA}^{(s)}\}_{s=1}^M$ attached after each vector self-attention block. Only the adapter corresponding to the active domain processes the flow descriptors; the geometric backbone remains shared.

Let x be the input features for a sample with domain s . Adapter $\text{FCA}^{(s)}$ generates a per-channel scale and bias from its flow vector:

$$\begin{aligned} \sigma^{(s)}, \mu^{(s)} &= \text{MLP}^{(s)}(C^{(s)}), \\ y^{(s)} &= \mathcal{P}_{\text{out}}^{(s)}\left(\left(\mathcal{P}_{\text{in}}^{(s)}(x) + \mu^{(s)}\right) \odot \sigma^{(s)}\right), \\ x &\leftarrow x + y^{(s)}. \end{aligned} \quad (4)$$

Here $\mathcal{P}_{\text{in}}^{(s)}$ and $\mathcal{P}_{\text{out}}^{(s)}$ are projections composed of Linear $\rightarrow$ LayerNorm $\rightarrow$ GELU. For mini-batches containing mixed domains, all adapters are computed in parallel and a one-hot routing mask selects the appropriate $y^{(s)}$ per sample, enabling efficient parallel training.

3.4. Training Objective and Inference

Given ground-truth pressure coefficients $\hat{P}$, the model minimizes an ℓ_1 regression loss on the predictions P . At inference time, selecting the adapter by the domain identifier s allows the same backbone to process data from cars, trains, aircrafts, or other scenarios, each under its native flow parameterization, while preserving a unified representation across different data sources.

4. Experiments

The scarcity of aerodynamic data is caused by multiple reasons: data generation requires a large amount of computational effort, much of the data is often private, and data from different subfields are usually considered incompatible. This leads to a lack of sufficient data in many subfields to train general models with strong generalization capabilities, greatly limiting the effectiveness of neural network methods. Therefore, we propose to combine flow field data from different subfields to make up for the insufficiency of data in individual subfields and train a flow field representation model with domain generality. In this section, we introduce a total of five sub-datasets from different fields, including an automotive dataset, a train dataset, an aircraft wing dataset, a full aircraft dataset, and a general shape dataset. We jointly train UniField on all datasets and verify its performance via multiple experiments:

1) We compare our model with existing models on the public automotive dataset DrivAerNet++ and prove that our model significantly reduces the prediction error of the surface pressure field compared with other models.

2) To verify the performance advantage brought by joint training, we train UniField on each dataset separately with the same model scale and compare it with the jointly trained model. The results show that the performance of the jointly trained model is consistently better than that of the single-dataset trained models. Especially in fields with relatively scarce data, such as the train and aircraft datasets, the error of the jointly trained model is more than one-third lower than that of the single-dataset models.

3) Considering that our method is mainly targeted at the aerodynamic subfield with scarce data, we investigated the impact of data volume on model performance. We conducted ablation experiments on the amount of data used for the aircraft wing dataset. Additionally, we explored the scenario of training with less aircraft wing data while keeping the settings of other datasets unchanged. In this case, we observed that the jointly trained model demonstrated a greater performance advantage.

4) We study the impact of model scale. We investigate three scales of models, namely 250M, 1B, and 2B. The experimental results show that the error of the model on each dataset decreases significantly with the increase in model

Model	DrivAerNet++			
	$\downarrow$ MSE ($\times 10^{-2}$)	$\downarrow$ MAE ($\times 10^{-1}$)	$\downarrow$ ReL2 (%)	$\downarrow$ ReL1 (%)
RegDGCNN [3]	8.29	1.61	27.72	26.21
Transolver [18]	7.15	1.41	23.87	22.57
FigConvNet [2]	4.99	1.22	20.86	21.12
TripNet [1]	5.14	1.25	20.05	20.93
UniField (ours)	4.53	1.00	18.67	16.22

Table 1. Model performance comparison on DrivAerNet++.

scale. This indicates that scaling neural networks have the potential to learn cross-domain universal representations of surface pressure fields.

4.1. Datasets

Pressure data standardization: To integrate data from different fields, we uniformly use the pressure coefficient (the dimensionless result of pressure) as the prediction target. In Fig.3, we present three samples from each dataset, including geometry and pressure coefficients. During testing, for DrivAerNet++, to facilitate comparison with other models, we follow the official default settings: we use the original pressure value subtracting the mean (-94.5) and dividing by the standard deviation (117.25) as the ground truth for evaluation (and also perform the corresponding transformation on the model output results). For other datasets other than DrivAerNet++, we continue to use the pressure coefficient (the dimensionless result of pressure) as the ground truth.

DrivAerNet++ [5]: DrivAerNet++ is a publicly available large-scale, high-fidelity dataset tailored for learning-based aerodynamic analysis of vehicles. All data are generated from high-resolution CFD simulations using steady-state incompressible flow solvers. Each sample captures the 3D surface pressure field of a vehicle traveling at 30 m/s. In total, DrivAerNet++ contains over 8,000 samples of car model surface point clouds with corresponding pressure fields, each consisting of more than 500,000 surface points. This rich dataset supports various aerodynamic prediction tasks, including lift and drag estimation, as well as dense parameter predictions such as surface pressure and wall shear stress. Our focus in this work is primarily on the surface pressure prediction task. The train-test split follows the official setting of DrivAerNet++¹.

Train dataset [11]: The train dataset comprises geometries of one high-speed train and two maglev trains, with CFD simulations conducted under multiple flow conditions. Specifically, flow conditions are characterized by the train’s traveling speed and the lateral wind speed perpendicular to

¹<https://github.com/Mohamedelrefaie/DrivAerNet/>




Figure 3. Visualization of the geometries and pressure distributions in each sub-datasets.

Dataset	Training methods		$\downarrow$ MSE ($\times 10^{-2}$)	$\downarrow$ MAE ($\times 10^{-1}$)	$\downarrow$ Rel L2 (%)	$\downarrow$ Rel L1 (%)
	Joint-Training	Single-Dataset				
DrivAerNet++		✓	4.47	1.04	19.61	16.88
	✓		4.53	1.00	18.67	16.22
AircraftWing		✓	2.73	0.70	35.1	23.4
	✓		1.85	0.50	28.6	17.1
FullAircraft		✓	0.42	0.33	28.3	22.6
	✓		0.21	0.23	19.6	15.7
Train		✓	2.97	0.45	175.0	112.2
	✓		0.89	0.29	94.1	69.7
FlowBench		✓	0.06	0.11	59.7	63.8
	✓		0.07	0.10	60.4	58.5

Table 2. Comparison between joint-training and single-dataset training across different aerodynamic domains. Joint training consistently improves prediction accuracy on datasets with limited samples or complex flow characteristics (e.g., *Train*, *FullAircraft*, *AircraftWing*), while large and well-sampled datasets such as *DrivAerNet++* show comparable results under both settings. This demonstrates that multi-domain joint training effectively mitigates data scarcity and enhances generalization for underrepresented flow regimes.

the travel direction. For the high-speed train, speeds range from 270 km/h to 350 km/h in 10 km/h increments, and lateral wind speeds range from 0 m/s to 15 m/s in 5 m/s increments, yielding 36 samples. For each maglev train geometry, speeds vary from 70 m/s to 100 m/s in 10 m/s increments, and lateral wind speeds range from 0 m/s to 20 m/s in 5 m/s increments, producing 20 samples per geometry. The high-speed train and one of the maglev train data are used for training, while the rest maglev train is used for evaluation.

Aircraft wing dataset [17]: This dataset includes 50 aircraft wing geometries. Flow conditions are described by Mach number and angle of attack. For each geometry, Mach numbers span from 0.2 to 0.8 in 0.2 increments, and angles of attack vary from -15° to 15° in 5° increments, result-

ing in 28 samples per geometry, totally 1400 samples. For training and evaluating UniField, we split the datasets into 45 geometries as the training set and the rest five as the test set.

Full Aircraft dataset [17]: Our complete aircraft dataset contains the geometries of 100 types of aircraft. Similar to the aircraft wing dataset, Mach number and angle of attack are used as flow conditions. For each geometry, the Mach number varies from 0.4 to 2.4 in increments of 0.4, and the angle of attack varies from -3° to 12° in increments of 3° , resulting in 36 samples for each geometry and 3600 samples in total. We use 80% of the data (i.e., 80 geometries) as the training set and the rest for evaluation.

FlowBench: We adopt FlowBench [14] as the general shape dataset. Unlike the previous several transportation system




Figure 4. Surface pressure field prediction visualization. Each column corresponds to one dataset. The results of FlowBench are scaled by 10 for clearer visualization. The model successfully captures the main aerodynamic features across diverse domains, where the errors in most areas are close to zero.

Data Volume	Joint-Training	$\downarrow$ MSE ($\times 10^{-2}$)	$\downarrow$ MAE ($\times 10^{-1}$)
1260 (90%)	✓	2.73 1.85 ($\downarrow 32\%$)	0.70 0.50 ($\downarrow 29\%$)
140 (10%)	✓	16.63 5.52 ($\downarrow 67\%$)	2.18 1.09 ($\downarrow 50\%$)

Table 3. **Joint-training under data-scarcity situations.** Results on the *AircraftWing* dataset show that as the available training data decrease from 90% to 10% of the total, the performance gain of joint training becomes significantly larger. This highlights that joint multi-domain training is especially advantageous in data-scarce scenarios, where shared aerodynamic priors compensate for limited samples.

datasets, the 3D part of FlowBench describes the flow in a lid-driven cavity, that is, the flow caused by moving the top cover of a square cavity when an object is placed inside. With completely different shapes and flow types, it is used to verify the universality of UniField in flow field modeling. Considering that the original data of FlowBench is in voxel form, we extract the voxels at the edge of the object in the cubic cavity to approximate surface pressure distributions. For FlowBench, we use the Reynolds number as the flow condition. FlowBench contains 1,000 sets of 3D data. We use 900 of them for training and the remaining 100 for testing.

4.2. Metrics

We adopt the evaluation metrics used in the DrivAerNet++ leaderboard for a comprehensive assessment of model performance. These metrics include Mean Squared Error

(MSE), Mean Absolute Error (MAE), Relative L2 error (RelL2), and Relative L1 error (RelL1). **Lower is better** for all the metrics.

4.3. Benchmark Comparison on DrivAerNet++

Setup: Previous research results [1, 2] typically randomly select 32,768 points from point clouds for training and testing. In our study, for the 250M model, the model takes 32,768 points as input; for larger models (1B and 2B), due to memory limitations, we reduce the number of input points to 8,192. However, during testing, we still use 32,768 points. For the 2B and 7B models, we randomly divide these 32,768 points into 4 groups, each containing 8,192 points, and predict the pressure field through the model for each group, and finally concatenate them to form the complete output. In the comparison of this section, we use the results of the 2B model.

Results: Results are summarized in Table 1. UniField consistently outperforms all other benchmarks. Compared to other benchmarks, UniField significantly reduces all types of errors, with MSE being approximately 10% lower than the previous best model FigConvNet. In terms of mean absolute error (MAE), UniField is 20% lower than TripNet. These results establish UniField as the new SOTA in surface pressure prediction.

4.4. Joint-Training vs. Single Dataset Training

Setup: To verify the benefits brought by the joint training strategy, we conducted a set of comparative experiments between the joint training models and the single-dataset models. Specifically, for each dataset, we trained a model from scratch using only that dataset and compared its performance with that of the joint training model on the same

Dataset	Model Scale	$\downarrow$ MSE ($\times 10^{-2}$)	$\downarrow$ MAE ($\times 10^{-1}$)	$\downarrow$ Rel L2 (%)	$\downarrow$ Rel L1 (%)
DrivAerNet++	250M	5.22	1.11	20.42	18.09
	1B	4.95	1.09	19.97	17.86
	2B	4.53	1.00	18.67	16.22
AircraftWing	250M	2.15	0.62	32.7	21.0
	1B	2.08	0.57	30.3	19.1
	2B	1.85	0.50	28.6	17.1
FullAircraft	250M	0.24	0.26	22.4	17.7
	1B	0.22	0.25	21.2	16.7
	2B	0.21	0.23	19.6	15.7
Train	250M	0.95	0.31	96.5	72.9
	1B	0.99	0.32	99.8	76.5
	2B	0.89	0.29	94.1	69.7
FlowBench	250M	0.070	0.110	62.8	61.7
	1B	0.067	0.111	61.7	62.9
	2B	0.067	0.105	60.4	58.5

Table 4. **Performance comparison of UniField with different parameter scales across five aerodynamic datasets.** As the model size increases from 250M to 2B parameters, the prediction error decreases for various datasets and metrics, demonstrating the scaling effect of larger models in capturing cross-domain aerodynamic representations. Best results in each dataset are highlighted in **bold**.

dataset.

Results: Table 2 compares the performance of *joint training* and *single-dataset training* across five aerodynamic datasets. A clear trend can be observed: datasets with limited data or complex flow structures benefit more from joint training, whereas large datasets with relatively simple flow patterns show only marginal differences.

For instance, the *DrivAerNet++* dataset, which contains abundant automotive samples, achieves nearly identical accuracy under both settings (MSE 4.47 vs. 4.53×10^{-2}), indicating that single-domain data are sufficient for robust learning. In contrast, the *Train* dataset, which are relatively small in scale, exhibit significant improvements when jointly trained with other domains: the MSE drops from 2.97×10^{-2} to 0.89×10^{-2} . Similarly, *AircraftWing* and *FullAircraft* also benefits notably from joint training (MSE $2.73 \rightarrow 1.85 \times 10^{-2}$ for *AircraftWing* and $0.42 \rightarrow 0.21 \times 10^{-2}$ for *FullAircraft*), indicating that the complex aerodynamic interactions around the wing are better modeled when auxiliary data from other geometries are introduced. These results verify that joint multi-domain training effectively mitigates data scarcity and improves generalization in complex or underrepresented flow regimes, supporting the conclusion that cross-domain knowledge sharing is especially valuable in data-scarce aerodynamic contexts.

In addition, in Fig.4, we present the predicted surface pressure field results of the jointly trained model on the test data of each subfield. UniField can provide prediction results that are very close to the ground truth, with errors in the vast majority of regions approaching zero.

4.5. Analysis of Joint Training and Data Scarcity

Setup: One of the core purposes of our research is to propose flow field models that can be applied in areas where data is scarce. Therefore, we conducted a study on the volume of domain data. Specifically, for the aircraft wing dataset, we investigated the comparison between another set of single-dataset models and the joint training models, which were trained using only 10% of the aircraft wing data, while the settings for the other datasets remained unchanged. We verify whether the joint training model can perform better with less data by comparing the gap between the single dataset model and the joint training model in two scenarios.

Results: As shown in Table 3, the advantage of joint training becomes more pronounced as the data volume decreases. Under the full-data setting, the jointly trained model achieves slightly lower errors (MSE 1.85 vs. 2.73×10^{-2} , MAE 0.50 vs. 0.70×10^{-1}). However, when the available data are reduced to only 10%, the performance

gap widens substantially: the MSE of the joint-training model is merely one-third of that from single-dataset training (5.52 vs. 16.63×10^{-2}), and the MAE is approximately half (1.09 vs. 2.18×10^{-1}). These results demonstrate that joint training enables effective knowledge transfer, particularly in data-scarce regimes.

4.6. Analysis of Model Scale

Setup: To investigate the influence of model capacity on aerodynamic field prediction, we evaluate three UniField variants with different parameter scales (**250M**, **1B**, and **2B**) across five datasets, including DrivAerNet++, AircraftWing, FullAircraft, Train, and FlowBench. All models share the same architecture and training strategy, with the only difference being the number of layers and hidden dimensions. Each model is trained under identical flow condition distributions and optimization settings to ensure a fair comparison.

Results: As summarized in Table 4, larger models consistently achieve lower errors across all metrics (MSE, MAE, Rel L2, and Rel L1), revealing a performance scaling trend. Specifically, the 2B model achieves an average reduction of about 10% in MSE compared with the 250M model, showing its enhanced capability to learn fine-grained flow structures and generalize across domains. The improvement is particularly evident in the FullAircraft and DrivAerNet++ datasets, where the complex 3D geometry and flow interaction benefit from the increased representational capacity of the larger model. These results confirm that scaling up the UniField model strengthens its ability to extract shared aerodynamic priors across heterogeneous flow domains.

5. Conclusion

In this work, we presented **UniField**, a universal framework for surface pressure field prediction across heterogeneous aerodynamic domains. Unlike conventional domain-specific approaches that rely on abundant single-domain data, UniField enables joint multi-domain training to learn shared flow representations from diverse sources such as automobiles, trains, aircraft, and generic shapes. By combining a domain-agnostic geometric backbone with domain-specific flow-conditioned adapters, UniField effectively disentangles geometry and flow information, achieving both generality and specialization within a unified architecture.

Extensive experiments demonstrate three key findings. First, UniField establishes new state-of-the-art performance on the public DrivAerNet++ benchmark, verifying that multi-domain integration does not compromise single-domain accuracy. Second, joint training consistently outperforms single-dataset training, with the largest performance gains observed in data-scarce domains such as aircraft and train flows—confirming that cross-domain knowledge sharing alleviates data scarcity. Third, a steady im-

provement of model performance is witnessed as the model scale increases, indicating that UniField has the ability to extract common aerodynamic representations across different fluid subfields by expanding the model’s scale.

Overall, UniField provides a step toward foundation models for aerodynamic analysis, capable of leveraging heterogeneous data sources to build robust, transferable field representations, outperforming model with single-domain data. Future work will focus on extending UniField beyond surface pressure prediction to full flow field prediction, paving the way for large-scale, unified modeling of complex fluid phenomena.

References

- [1] Qian Chen, Mohamed Elrefaie, Angela Dai, and Faez Ahmed. Tripnet: Learning large-scale high-fidelity 3d car aerodynamics with triplane networks, 2025. 2, 3, 5, 7
- [2] Chris Choy, Alexey Kamenev, Jean Kossaifi, Max Rietmann, Jan Kautz, and Kamyar Azizzadenesheli. Factorized implicit global convolution for automotive computational fluid dynamics prediction, 2025. 2, 3, 5, 7
- [3] Mohamed Elrefaie, Angela Dai, and Faez Ahmed. Drivaernet: A parametric car dataset for data-driven aerodynamic design and graph-based drag prediction. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, page V03AT03A019. American Society of Mechanical Engineers, 2024. 5
- [4] Mohamed Elrefaie, Angela Dai, and Faez Ahmed. Drivaernet: A parametric car dataset for data-driven aerodynamic design and prediction. *Journal of Mechanical Design*, 147 (4), 2025. 3
- [5] Mohamed Elrefaie, Florin Morar, Angela Dai, and Faez Ahmed. Drivaernet++: A large-scale multimodal car dataset with computational fluid dynamics simulations and deep learning benchmarks, 2025. 2, 5
- [6] Hongrui Gao, Tanghong Liu, Xiaodong Chen, Xiaoshuai Huo, Zhengwei Chen, Jie Zhang, and Boo Cheong Khoo. An accurate and efficient methodology on wind spectra relative to moving trains: field measurements of wind characteristics in complex terrains. *Stochastic Environmental Research and Risk Assessment*, pages 1–31, 2025. 1
- [7] Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023. 3
- [8] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020. 3
- [9] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in neural information processing systems*, 33:11525–11538, 2020. 4

- [10] Huakun Luo, Haixu Wu, Hang Zhou, Lanxiang Xing, Yichen Di, Jianmin Wang, and Mingsheng Long. Transolver++: An accurate neural solver for pdes on million-scale geometries, 2025. [3](#)
- [11] Wei Qiu, Zhenxu Sun, Junxiang Wang, Ye Bai, Shuanbao Yao, Dilong Guo, and Guowei Yang. A compressed sensing based framework for surface pressure field reconstruction from sparse measurement. *Physics of Fluids*, 37(4), 2025. [1](#), [2](#), [3](#), [5](#)
- [12] M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. [3](#)
- [13] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.
- [14] Ronak Tali, Ali Rabeh, Cheng-Hau Yang, Mehdi Shadkhah, Samundra Karki, Abhisek Upadhyaya, Suriya Dhakshinamoorthy, Marjan Saadati, Soumik Sarkar, Adarsh Krishnamurthy, Chinmay Hegde, Aditya Balu, and Baskar Ganapathysubramanian. Flowbench: A large scale benchmark for flow simulation over complex geometries, 2024. [2](#), [6](#)
- [15] Alasdair Tran, Alexander Mathews, Lexing Xie, and Cheng Soon Ong. Factorized fourier neural operators. *arXiv preprint arXiv:2111.13802*, 2021. [3](#)
- [16] Tian Wang and Chuang Wang. Latent neural operator for solving forward and inverse pde problems. *Advances in Neural Information Processing Systems*, 37:33085–33107, 2024. [3](#)
- [17] Yueqing Wang, Peng Zhang, Yushuang Liu, Jianing Zhao, Jie Lin, and Yi Chen. Aerodynamic coefficients prediction via cross-attention fusion and physical-informed training. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(1):869–876, 2025. [2](#), [3](#), [6](#)
- [18] Haixu Wu, Huakun Luo, Haowen Wang, Jianmin Wang, and Mingsheng Long. Transolver: A fast transformer solver for pdes on general geometries, 2024. [3](#), [5](#)
- [19] Chang Yan, Shengfeng Xu, Zhenxu Sun, Thorsten Lutz, Dilong Guo, and Guowei Yang. A framework of data assimilation for wind flow fields by physics-informed neural networks. *Applied Energy*, 371:123719, 2024. [3](#)
- [20] Zhanhong Ye, Zining Liu, Bingyang Wu, Hongjie Jiang, Leheng Chen, Minyan Zhang, Xiang Huang, Qinghe Meng Zou, Hongsheng Liu, Bin Dong, et al. Pdeformer-2: A versatile foundation model for two-dimensional partial differential equations. *arXiv preprint arXiv:2507.15409*, 2025. [3](#)
- [21] Navid Zehtabiyani-Rezaie, Alexandros Iosifidis, and Mahdi Abkar. Data-driven fluid mechanics of wind farms: A review. *Journal of Renewable and Sustainable Energy*, 14(3): 032703, 2022. [2](#)
- [22] Hao Zhang, Yang Shen, Wei Huang, Zan Xie, and Yao-Bin Niu. Deep transfer learning for three-dimensional aerodynamic pressure prediction under data scarcity. *Theoretical and Applied Mechanics Letters*, 15(2):100571, 2025.
- [23] Qiao Zhang, Xuan Zhao, Kai Li, Xinwu Tang, Jifei Wu, and Weiwei Zhang. Prediction model of aircraft hinge moment: Compressed sensing based on proper orthogonal decomposition. *Physics of Fluids*, 36(7), 2024. [2](#)
- [24] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun. Point transformer, 2021. [2](#), [4](#)
- [25] Hang Zhou, Yuezhou Ma, Haixu Wu, Haowen Wang, and Mingsheng Long. Unisolver: Pde-conditional transformers are universal pde solvers. *arXiv preprint arXiv:2405.17527*, 2024. [3](#)
- [26] Junhong Zou, Wei Qiu, Zhenxu Sun, Xiaomei Zhang, Zhaoxiang Zhang, Xiangyu Zhu, and Zhen Lei. Joint optimization of sensor placement and sparse pressure field reconstruction with a two-stage framework for limited data. *Physics of Fluids*, 37(7), 2025. [3](#)