CAAI Transactions on Intelligence Technology

WILEY

The Institution of Engineering and Technology

ORIGINAL RESEARCH OPEN ACCESS

APTNet: A Condition-Sensitive Modulation Framework for Surface Pressure Prediction on Supersonic Aircraft

Hongbin Xu[1] | Yin Long[1] | Junlin Wu[2]

[1]School of Computer Science and Technology, Southwest University of Science and Technology, Mianyang, China | [2]National Key Laboratory of Aerospace Physics in Fluids, Mianyang, China

Correspondence: Yin Long (yinlong@swust.edu.cn)

Received: 10 January 2026 | **Revised:** 27 April 2026 | **Accepted:** 12 May 2026

Keywords: aerodynamic prediction | condition-sensitive local modulation | flow-aligned graph attention | physics-informed deep learning | surface pressure prediction

ABSTRACT

Accurate surface-pressure prediction over broad operating envelopes is critical for supersonic aerodynamic analysis and design. To overcome the bottlenecks of traditional computational fluid dynamics (CFD) in real-time performance and computational efficiency, data-driven deep learning methods have emerged. However, existing data-driven models face two fundamental physical challenges: first, the pressure response to Mach number, angle of attack and altitude is strongly spatially heterogeneous, which cannot be adequately captured by traditional global condition fusion methods; second, the inherent strong anisotropy of supersonic flows renders standard isotropic graph construction based on Euclidean distance prone to spurious cross-shock connections, resulting in nonphysical smoothing of predictions. To address these issues, we propose aircraft pressure transformer network (APTNet), a physics-informed deep learning framework that explicitly encodes flow–geometry interactions and streamwise physical priors. Specifically, we first propose a condition-sensitive local modulation (CSLM) module, which utilises a surface-normal–freestream alignment term to distinguish windward and leeward regions, and learns per-point sensitivity fields to implement physically interpretable local pressure modulation via a feature-wise linear modulation (FiLM) mechanism. Furthermore, we introduce a flow-aligned anisotropic graph construction strategy, which reshapes the local receptive fields of graph nodes into ellipsoids elongated along streamlines, enforcing information to propagate preferentially along streamlines and suppressing cross-shock feature mixing. Experiments on diverse complex aircraft configurations demonstrate that APTNet consistently outperforms existing baseline models in terms of both accuracy and efficiency. Beyond quantitative improvements, the model provides interpretable intermediate signals, such as attention maps and feature responses, which can be directly integrated into the aerodynamic design workflow.

1 | Introduction

Accurate prediction of surface pressure distributions across varying flow conditions is fundamental to aerodynamic analysis, structural load assessment and design optimisation of supersonic and transonic aircraft. Although high-fidelity CFD offers reliable solutions, the associated computational cost becomes prohibitive for rapid design iteration, uncertainty quantification or near real-time evaluation. These limitations have motivated

Abbreviations: AoA, angle of attack; APTNet, aircraft pressure transformer network; CFD, computational fluid dynamics; CSLM, condition-sensitive local modulation; FAGAT, flow-aligned graph attention layer; FiLM, feature-wise linear modulation; GNNs, graph neural networks; k-NN, k-nearest neighbour; Ma, Mach number; MLP, multilayer perceptron; MRE, mean relative error; NS, Navier–Stokes; RANS, Reynolds-averaged Navier–Stokes; Re, Reynolds number; RF, receptive field.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2026 The Author(s). *CAAI Transactions on Intelligence Technology* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology and Chongqing University of Technology.

the development of data-driven surrogate models, which provide substantially faster inference whilst retaining a large portion of the predictive capability of CFD [1, 2].

Beyond forward surrogate prediction, intelligent aerodynamics has recently extended towards inverse design and shape optimisation with generative models. In particular, diffusion-based and adversarial frameworks have shown promising capability in high-dimensional aerodynamic design tasks, including three-dimensional aircraft inverse design, stealth-oriented shape optimisation and multipoint inverse design of flying-wing configurations [3–5]. These developments suggest that intelligent aerodynamics is evolving from pure flowfield prediction towards a broader paradigm integrating fast surrogate modelling, inverse design and optimisation. Therefore, accurate and efficient forward prediction models remain an important foundation for downstream design-oriented aerodynamic intelligence, especially when physically complex quantities, such as surface pressure distributions over three-dimensional high-speed configurations, are involved.

However, learning an accurate surrogate for surface pressure remains challenging due to the strongly spatially heterogeneous nature of high-speed aerodynamic flows. Changes in Mach number modify compressibility and shock intensity; variations in angle of attack (AoA) induce asymmetric windward-leeward loading and pressure-centre shifts and altitude affects atmospheric density and dynamic pressure in a nonuniform manner across the surface. Crucially, these influences manifest locally: stagnation regions, leeward surfaces, shock-prone areas and high-curvature features respond differently to identical condition variations. Modelling this localised, condition-dependent sensitivity is essential for accurate pressure prediction and robust generalisation.

Recent advances in deep learning for fluid mechanics have shown the effectiveness of graph neural networks (GNNs) and mesh-based learning frameworks for capturing geometric complexity and spatial interactions. MeshGraphNet [6] demonstrated that message-passing networks can learn physical dynamics on unstructured meshes, inspiring a series of GNN-based aerodynamic surrogates [7]. Meanwhile, transformer architectures have been adapted to fluid prediction due to their strong long-range dependency modelling [8]. Despite these advancements, most existing surrogate models rely on global fusion mechanisms to incorporate flow-condition information, typically through direct concatenation [1] or token-based cross-attention to couple condition embeddings with geometric representations [9]. Such global fusion implicitly assumes that flow conditions exert a uniform influence across the entire surface. Although this is often sufficient for predicting integrated aerodynamic coefficients, it is inadequate for surface pressure fields, whose sensitivity to Mach, AoA and altitude varies drastically across different geometric regions. Consequently, global fusion approaches struggle to capture orientation-dependent behaviours, shock-sensitive responses and subtle redistribution patterns that are critical in high-speed aerodynamics.

A second, often overlooked, challenge lies in how the neighbourhood graph is constructed for unstructured surface learning. Most existing methods rely on a standard k-nearest-neighbour (k-NN) algorithm based on the Euclidean distance to build the underlying topology. Such an isotropic neighbourhood search mechanism implicitly assumes that a target point is equally correlated with its nearby points in any spatial direction. However, this assumption contradicts the physical nature of supersonic aerodynamics, where the flow field exhibits strong anisotropy: physical quantities (e.g., pressure and density) evolve relatively smoothly and continuously along streamlines, whilst they can change abruptly or even discontinuously in directions transverse to the flow, especially in the presence of strong-gradient structures such as shock waves or expansion fans. With an isotropic k-NN graph, the receptive field is essentially spherical and cannot distinguish flow-aligned from cross-flow correlations. As a result, message passing can easily bridge shock interfaces, mistakenly aggregating features from the preshock (low-pressure) and postshock (high-pressure) regions, which leads to nonphysical smoothing near shock fronts and substantially weakens the model's ability to resolve high-frequency pressure gradients. To break this limitation, it is necessary to explicitly incorporate flow-directional physical priors into the low-level graph construction. This motivates a flow-aligned anisotropic graph construction strategy that reshapes local receptive fields to favour along-streamline aggregation whilst suppressing cross-shock connections, thereby alleviating shock blurring and enabling a more faithful reconstruction of complex supersonic flow structures.

To address these limitations, we propose APTNet, a physics-informed deep learning framework that explicitly models the interaction between flow conditions and local geometric features. The core of APTNet is the CSLM module, which constructs a per-point sensitivity field to characterise how each surface point responds differently to changes in Mach number, AoA and altitude. CSLM incorporates a flow–geometry alignment term based on the dot product between surface normals and the freestream direction, enabling the model to distinguish windward, leeward and stagnation-dominated regions. Through a FiLM-style scale-and-shift transformation, CSLM injects localised physically interpretable condition-dependent effects into the geometric features. These condition-modulated features are further processed by graph attention and multiscale transformer blocks to capture both local topological interactions and global aerodynamic coupling. Specifically, the main contributions of this work are summarised as follows:

- We propose a CSLM mechanism that embeds flow–geometry interaction via a surface-normal–freestream alignment indicator and learns per-point sensitivity fields for Mach, angle of attack and altitude. These sensitivities are injected using FiLM, enabling localised pressure amplification and redistribution with physical interpretability.
- We propose an anisotropic graph-construction strategy aligned with flow direction. By replacing isotropic Euclidean distance with a weighted flow-aware metric, the node neighbourhood becomes an ellipsoid stretched along streamlines encouraging message passing along the flow. This suppresses nonphysical smoothing across shock waves and improves the accuracy and sharpness of shock-front capture.
- We design a surface-based learning pipeline that combines topology-aware local propagation with configuration-level

coupling, improving pressure prediction accuracy and cross-geometry generalisation under varying flow conditions, whilst yielding interpretable sensitivity maps consistent with aerodynamic intuition (e.g., Mach-dominated stagnation regions and AoA-dominated windward–leeward regions).

2 | Related Work

2.1 | Deep-Learning Surrogate Models for Aerodynamics

Data-driven surrogate models have been widely explored to accelerate CFD by learning mappings from geometry and flow conditions to pressure or velocity fields [10, 11]. For two-dimensional aerofoils, Thuerey et al. demonstrated that U-net-based convolutional networks can efficiently infer Reynolds-averaged Navier–Stokes (RANS) solutions. Related CNN-based aerofoil surrogates for rapid pressure/flow prediction have also been reported in the literature [12]. More broadly, Vinuesa and Brunton reviewed how deep learning can enhance CFD, emphasising its potential to accelerate simulations, improve turbulence modelling and build reduced-order models for engineering flows [2, 13, 14].

For three-dimensional configurations, Shen et al. proposed a deep-learning framework based on PointNet++ to predict surface pressure coefficients on general flight vehicles, including slender bodies, spaceplanes and rockets, using point-cloud representations instead of volume grids [15]. Their results highlight that point-based neural networks can generalise across families of 3D geometries with varying flow conditions. Complementary point-cloud-based aerodynamic prediction efforts have also been explored for 3D wings/aircraft or vehicle-relevant geometries [16–19]. More recent work has introduced attention mechanisms into aerodynamic surrogates, such as attention-enhanced U-net architectures for transonic aerofoil RANS inference with improved robustness to shock-induced separation [20] and deep attention networks for fast aerodynamics prediction of laminar aerofoils [21].

These studies collectively show that deep neural networks can provide accurate and efficient aerodynamic surrogates on both 2D and 3D configurations. However, most existing models operate on regular grids or global point sets and typically treat flow conditions as globally shared inputs, without explicitly representing the spatially nonuniform sensitivity of surface pressure to Mach number, angle of attack and altitude across complex aircraft geometries. This motivates architectures, such as APTNet, which couple local condition-sensitive modelling with mesh-aware propagation.

2.2 | GNNs for Aerodynamic Surrogates

GNNs have emerged as a natural choice for learning aerodynamic fields defined on irregular meshes or point clouds [22, 23]. Their ability to operate on non-Euclidean domains makes them particularly suitable for external aerodynamic configurations, where surface geometry is typically represented by unstructured triangular meshes. Early advances, such as MeshGraphNet [6], demonstrated that message-passing frameworks can accurately propagate physical information across mesh connectivity and capture complex flow interactions. Building on this idea, GNN-based surrogate models have been explored to accelerate CFD simulations by learning mesh-based flow evolution [24, 25].

More recently, several studies have applied GNNs to aerodynamic prediction tasks directly relevant to the present work. Hines and Bekemeyer [26] used a graph-based architecture to predict surface pressure distributions on both aerofoils and full aircraft geometries, showing that GNNs can effectively encode geometric neighbourhood information and improve prediction quality relative to point-based models. In the context of vehicle aerodynamics, Roznowicz et al. [7] developed large-scale graph-machine-learning surrogates for three-dimensional external flowfields around complex automotive shapes, highlighting the scalability of graph-based approaches and their ability to model multibody interactions and wake structures. Large-scale 3D aerodynamic datasets and learning pipelines further support the practical relevance of mesh-graph-based surrogates in design workflows [18].

These studies collectively demonstrate that GNNs provide an effective mechanism for encoding surface topology and enforcing local geometric consistency, attributes that are crucial for modelling the spatially varying pressure distribution on aircraft surfaces. However, existing GNN-based aerodynamic surrogates typically integrate flow conditions through global concatenation or node-wise conditioning, without explicitly modelling the spatially nonuniform sensitivity of pressure to Mach number, AoA or altitude. This motivates the development of our CSLM module, which introduces condition-aware per-point modulation prior to graph propagation, enabling the network to represent aerodynamic responses that depend on both local geometry and flow–geometry alignment.

2.3 | Transformer and Attention-Based Architectures for Flow Prediction

Several works have proposed attention-augmented surrogates that combine convolutional encoders with self-attention to better capture long-range spatial dependencies in flow fields. For example, attention U-net variants have been developed to infer transonic aerofoil RANS solutions on structured grids, improving robustness in regions with shocks and large-scale separation [20]. A physics-integrated attention model has also been introduced for rapid two-dimensional flow prediction, where flow channel geometry is embedded as images and processed with transformer-style encoders enriched by physical information in the loss function [27]. Beyond grid surrogates, transformer-family models have also been used for nonintrusive flow inference and reduced-order modelling, demonstrating their ability to capture nonlocal dependencies in fluid dynamics [8, 28].

More directly, transformer-based decoders have been employed to accelerate numerical simulations. A recent Physics of Fluids study proposed a transformer-based deep-learning solver for incompressible flows on coarse grids, where self-attention replaces part of the traditional numerical discretisation whilst maintaining accuracy on benchmark problems [29]. In parallel,

physics-informed transformer variants have been explored to embed governing equations or physical constraints into the attention framework, improving stability and generalisation for flow forecasting and super-resolution [30, 31]. These works demonstrate that attention mechanisms are effective for modelling nonlocal interactions and multiscale structures in fluid dynamics.

Nevertheless, most existing transformer-based flow surrogates operate on Eulerian grid fields or image-like representations, and typically do not explicitly incorporate surface mesh topology or per-point geometric features for complex aircraft configurations. Moreover, flow conditions are usually encoded as global tokens or concatenated parameters rather than driving spatially varying modulation at the surface level. In contrast, APTNet combines mesh-based operators with attention-based global coupling and introduces CSLM to realise condition-sensitive local modulation on the aircraft surface.

2.4 | Shape–Condition Fusion Strategies in Aerodynamic Surrogates

A central design choice in aerodynamic surrogate models concerns how to fuse geometric information with flow conditions. Many early deep-learning methods represent geometry on regular grids (e.g., signed-distance or occupancy fields) and concatenate flow parameters, such as Mach number and angle of attack, as additional input channels or global conditioning vectors. This strategy is widely used in CNN-based surrogates for aerofoil and wing flows, including U-net models and attention-based networks for laminar aerofoil aerodynamics [21]. For 3D vehicles, PointNet++-based frameworks often encode the point-cloud geometry and simply append flow conditions to the global feature vector or to each point feature before the final regression layers [15].

Recent attention-based models follow a similar pattern: geometric features are first extracted from images or grids, and flow conditions are injected either via global tokens or concatenated condition vectors before the attention blocks [20, 27]. Although such global fusion schemes are simple and effective, they inherently assume that all points on the surface respond in a similar way to changes in flow conditions.

For complex aircraft geometries in transonic and supersonic regimes, the surface pressure sensitivity to Mach number, AoA and altitude is strongly nonuniform: stagnation regions, leading edges and different oriented surface panels exhibit qualitatively different responses under the same condition change. Global fusion or node-wise conditioning without explicit flow–geometry alignment cannot directly encode this heterogeneity, and the network must implicitly infer it through deeper layers. The CSLM module in APTNet is designed to address this limitation: it constructs a per-point condition-sensitivity field based on geometric features, surface-freestream alignment and independently encoded flow conditions, and then uses a FiLM-style modulation [32, 33] to inject these sensitivities into the local features before graph and transformer propagation.

3 | APTNet Framework

3.1 | Overview of APTNet

APTNet is an end-to-end physics-informed surrogate that maps an aircraft's three-dimensional surface geometry and flow conditions to its surface-pressure distribution. The architecture decomposes pressure formation into three physically motivated stages followed by a prediction head: (i) geometry encoding at each surface point, (ii) CSLM that injects Mach/AoA/altitude effects in a spatially nonuniform way and (iii) topology-/range-aware propagation (a flow-aligned graph attention layer for local neighbourhood coupling and a multiscale transformer for configuration-wide interactions) before regressing pressure. An overview of the framework is shown in Figure 1.

3.1.1 | *Physical Motivation*

In transonic/supersonic regimes, surface pressure exhibits spatially heterogeneous sensitivity to flow conditions: stagnation-oriented, windward and leeward regions respond differently to changes in Mach number, angle of attack and altitude. Meanwhile, the pressure field is globally constrained and couples across the surface through both mesh topology (short-range) and aerodynamic interactions (long-range). APTNet reflects these properties by first building per-point condition-aware features (CSLM) and then propagating them with locality- and range-sensitive operators (a flow-aligned graph attention layer and a transformer).

3.1.2 | *Coordinate Preprocessing*

Given surface points $X = \{x^{(i)}\}_{i=1}^{N}$, $x^{(i)} \in \mathbb{R}^3$ (tensor shape $[B, N, 3]$), we apply sample-wise decentralisation and scale normalisation prior to point encoding and graph construction:

$$\bar{x} = \frac{1}{N}\sum_{i=1}^{N} x^{(i)}, \quad \tilde{x}^{(i)} = x^{(i)} - \bar{x},$$
$$s = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \|\tilde{x}^{(i)}\|_2^2}, \quad \hat{x}^{(i)} = \tilde{x}^{(i)}/(s+\varepsilon). \tag{1}$$

This translation-/scale-invariant normalisation stabilises neighbourhood search and feature magnitudes whilst preserving local shape cues (e.g., curvature and edges) that govern pressure variation.

3.1.3 | *Overall Processing Pipeline of APTNet*

From normalised coordinates $\{\hat{x}^{(i)}\}$, we extract per-point geometric features with a lightweight point encoder, $\phi(\hat{x}^{(i)}) \in \mathbb{R}^H$. The CSLM module (Section 3.2) then constructs a per-point condition-sensitivity field by combining (a) independently encoded flow conditions (Mach, AoA and altitude), (b) a flow–geometry alignment term (unit normal dotted with freestream direction) and (c) a FiLM-style scale-shift transformation that injects condition effects into $\phi(\hat{x}^{(i)})$ to produce condition-modulated features $\tilde{h}^{(i)} \in \mathbb{R}^H$. Next, a flow-aligned graph

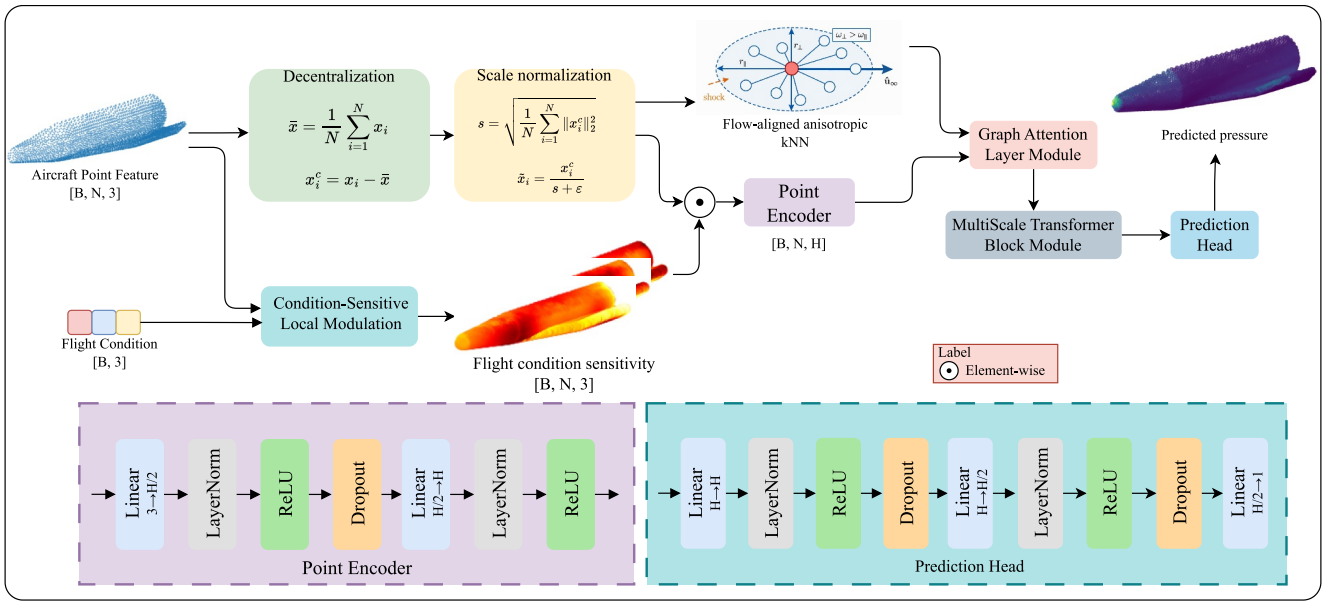


FIGURE 1 | APTNet framework.

attention layer (FAGAT) (Section 3.3) operates on a k-NN surface graph constructed under a flow-aligned weighted distance from $\{\hat{x}^{(i)}\}$, performing adjacency-aware aggregation to enforce local pressure consistency whilst suppressing spurious cross-shock feature mixing and capturing short-range aerodynamic coupling. Finally, a multiscale transformer (Section 3.4) applies stacked self-attention to $\{z^{(i)}\}$ (the FAGAT outputs), expanding the effective receptive field with depth to model long-range interactions (e.g., wing-body interference, shock-shock and shock-BL coupling) and yielding globally informed, condition-aware node representations $\{\hat{z}^{(i)}\}$. A lightweight MLP head maps $\{\hat{z}^{(i)}\}$ to surface pressure.

3.1.4 | *Physical Rationale for the Module Design*

CSLM resolves orientation-dependent region-specific responses that global fusion cannot represent; FAGAT enables flow-aligned topology-constrained neighbourhood propagation on the surface; the transformer captures nonlocal configuration-wide coupling. This division of roles mirrors the physical nature of the pressure field and underpins the empirical gains reported in Section 4.3.

3.2 | Condition-Sensitive Local Modulation

In supersonic and transonic aerodynamics, the surface pressure field exhibits strongly spatially heterogeneous responses to variations in Mach number, AoA and altitude. Such variations affect distinct aerodynamic mechanisms: Ma alters compressibility and shock intensity, AoA changes windward–leeward pressure redistribution and altitude modifies atmospheric density and dynamic pressure. Therefore, each surface point responds differently to flow-condition variations.

To capture this phenomenon, APTNet incorporates the CSLM module, whose overall framework is illustrated in Figure 2. ItFigure 2 constructs a per-point condition-sensitivity field and modulates geometric features using a physically interpretable FiLM-style transformation.

3.2.1 | *Flow-Condition Encoding*

Let $c_k \in \{Ma, AoA, Alt\}$ denote the k-th flow condition. Each condition is first encoded independently by a lightweight MLP

$$e_k = f_k(c_k), \qquad e_k \in \mathbb{R}^{H/3}, \tag{2}$$

where $f_k(\cdot)$ is a learnable function. The encoded condition vectors are then concatenated to form a unified condition representation

$$e_{\text{cat}} = [e_{Ma}, e_{AoA}, e_{Alt}] \in \mathbb{R}^{H}. \tag{3}$$

where Alt denotes altitude.

3.2.2 | *Flow–Geometry Alignment Term*

To quantify the orientation-dependent aerodynamic response of each surface point, we define an alignment scalar

$$d^{(i)} = \mathbf{n}^{(i)} \cdot \hat{\mathbf{u}}_{\infty}, \tag{4}$$

where $\mathbf{n}^{(i)} \in \mathbb{R}^3$ is the outward unit surface normal of point i and $\hat{\mathbf{u}}_{\infty}$ is the unit incoming-flow direction. For flows without sideslip, the freestream direction is uniquely determined by the angle of attack α as

$$\hat{\mathbf{u}}_{\infty} = \begin{bmatrix} \cos\alpha \\ 0 \\ \sin\alpha \end{bmatrix}, \tag{5}$$

which corresponds to the standard body-fixed aerodynamic coordinate system (x forward and z downward).

The alignment value satisfies $d^{(i)} \in [-1, 1]$ and characterises the relative exposure of the surface to the incoming flow. Points

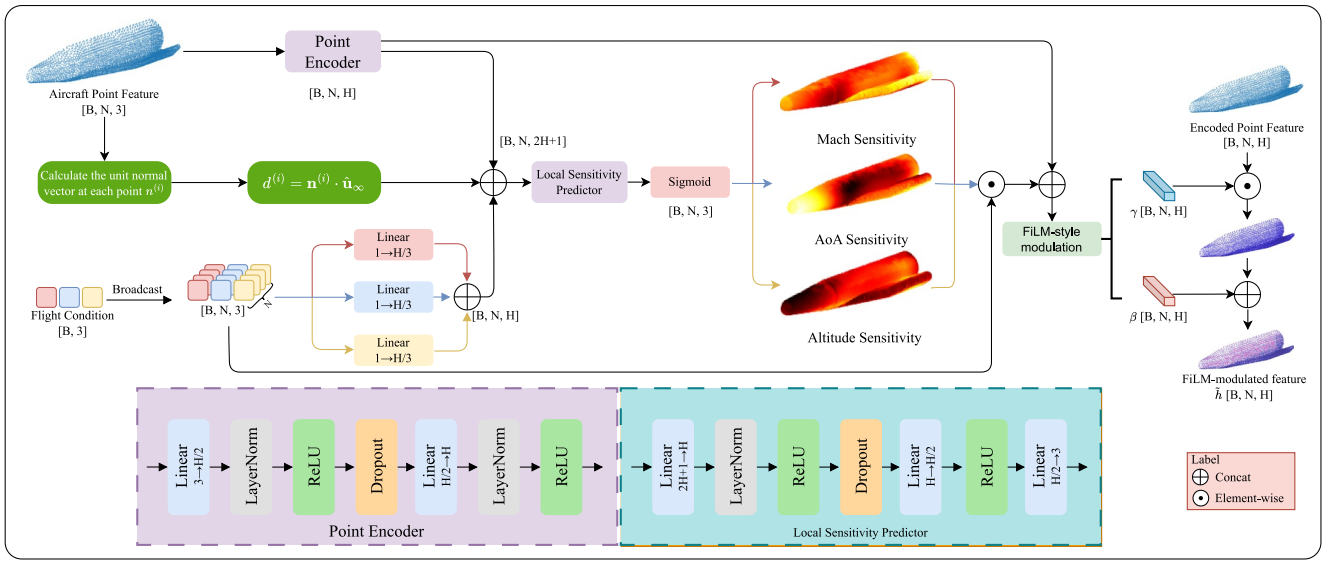


FIGURE 2 | CSLM framework.

with large $d^{(i)}$ are more directly facing the freestream and tend to correspond to windward or stagnation-oriented regions, where compressibility effects and shock-related pressure amplification are pronounced. Points with $d^{(i)} \approx 0$ lie on lateral surfaces, where the local pressure is closer to the freestream static level and the flow is mainly influenced by shear and boundary-layer development. Points with $d^{(i)} < 0$ correspond to leeward regions, where the flow undergoes expansion and may experience separation and wake formation, making the pressure more sensitive to changes in angle of attack.

Although $d^{(i)}$ is computed from simple geometric and flow-alignment information, it provides CSLM with a physically meaningful indicator of whether a surface point is in a 'compression-dominated' or 'expansion-dominated' state. This allows the subsequent sensitivity estimation to distinguish, from the outset, between windward, lateral and leeward regions rather than relying solely on deeper layers to infer such orientation-dependent behaviour from geometry alone.

3.2.3 | *Point-Wise Sensitivity Estimation*

Let $\phi(\hat{x}^{(i)}) \in \mathbb{R}^H$ be the geometric feature of point i obtained from the point encoder. We then concatenate the point feature, the flow–geometry alignment scalar and the unified condition representation to form

$$s^{(i)} = \left[\phi(\hat{x}^{(i)}),\ d^{(i)},\ e_{\text{cat}}\right] \in \mathbb{R}^{2H+1}. \tag{6}$$

A local sensitivity predictor $g(\cdot)$ is used to produce the sensitivities to Mach number, angle of attack and altitude jointly:

$$\left[w_{Ma}^{(i)},\ w_{AoA}^{(i)},\ w_{Alt}^{(i)}\right] = \sigma\left(g\left(s^{(i)}\right)\right), \tag{7}$$

where $g(\cdot)$ is a small MLP, $\sigma(\cdot)$ is the sigmoid function and each sensitivity weight satisfies

$$w_{Ma}^{(i)},\ w_{AoA}^{(i)},\ w_{Alt}^{(i)} \in (0, 1). \tag{8}$$

The collection $\{w_k^{(i)}\}_{i=1}^N$ forms a condition-specific attention-like sensitivity field over the surface. For example, during training the network can learn to assign larger $w_{\text{Ma}}^{(i)}$ in stagnation and leading-edge regions, where pressure is strongly affected by compressibility and shock strength; larger $w_{\text{AoA}}^{(i)}$ in windward and leeward regions, where pressure redistribution is primarily driven by changes in angle of attack; and larger $w_{Alt}^{(i)}$ in regions where the pressure level mainly follows dynamic-pressure scaling with altitude. Importantly, these patterns are not imposed a priori but are learnt from data, allowing CSLM to break the implicit assumption of spatially uniform condition sensitivity that underlies simple global fusion schemes.

3.2.4 | *Composite Condition Representation*

The scalar sensitivity weights are applied to the independently encoded condition vectors through scalar weighting followed by concatenation. Specifically, we first form a weighted concatenation:

$$e_{\text{cond}}^{(i)} = \left[w_{Ma}^{(i)} e_{Ma},\ w_{AoA}^{(i)} e_{AoA},\ w_{Alt}^{(i)} e_{Alt}\right] \in \mathbb{R}^H. \tag{9}$$

This formulation preserves the individual contributions of Mach number, angle of attack and altitude whilst yielding a unified condition representation with the same dimensionality as the geometric feature.

3.2.5 | *FiLM-Style Modulation*

After constructing the point-wise condition representation $e_{\text{cond}}^{(i)} \in \mathbb{R}^H$, CSLM injects flow-condition effects into the geometric features through a FiLM-style modulation. Specifically, we concatenate the geometric feature and the condition representation and compute a pair of scale and shift coefficients:

$$\left[\gamma^{(i)}, \beta^{(i)}\right] = h\left(\left[\phi(\hat{x}^{(i)}),\ e_{\text{cond}}^{(i)}\right]\right), \tag{10}$$

where $h(\cdot)$ is a lightweight MLP that maps $\mathbb{R}^{2H}$ to $\mathbb{R}^{2H}$ and $\gamma^{(i)}, \beta^{(i)} \in \mathbb{R}^H$. The final modulated feature is defined as

$$\tilde{h}^{(i)} = \gamma^{(i)} \odot \phi(\hat{x}^{(i)}) + \beta^{(i)}, \tag{11}$$

where $\odot$ denotes element-wise multiplication.

This FiLM transformation serves two key aerodynamic purposes. First, the scale term $\gamma^{(i)}$ modulates the magnitude of the geometric feature $\phi(\hat{x}^{(i)})$, enabling APTNet to capture Mach-dependent pressure amplification near stagnation regions and altitude-dependent dynamic-pressure scaling across the surface. Second, the shift term $\beta^{(i)}$ introduces localised feature offsets, representing AoA-driven pressure redistribution (e.g., increased loading on the windward side and relief on the leeward side). Together, the FiLM-modulated feature $\tilde{h}^{(i)}$ encodes how each surface point reacts to combined flow conditions, providing an aerodynamically consistent input for the subsequent graph attention and multiscale transformer blocks.

3.2.6 | *Condition-Modulated Features for Subsequent Propagation*

The modulated features $\left\{\tilde{h}^{(i)}\right\}_{i=1}^{N}$ (where N is the number of surface points) are passed into the graph-attention and multiscale transformer modules for topology-aware spatial propagation and global coupling. The spatially heterogeneous sensitivity patterns that CSLM aims to capture are illustrated schematically in Figure 3.

3.2.7 | *Physical Interpretation of CSLM*

From a physical viewpoint, CSLM can be regarded as a three-step mechanism that maps local geometry and flow conditions to a condition-aware feature representation. First, the flow–geometry alignment term $d^{(i)}$ labels each surface point according to its incidence with respect to the freestream, distinguishing windward, lateral and leeward orientations. Second, the point-wise sensitivity weights $w_k^{(i)}$ form condition-specific, attention-like fields over the surface, indicating where the pressure should be more responsive to changes in Mach number, angle of attack or altitude. In particular, the network can learn to assign higher Mach sensitivity in stagnation and leading-edge regions, higher AoA sensitivity on windward and leeward panels and higher altitude sensitivity in regions where the pressure primarily follows dynamic-pressure scaling.

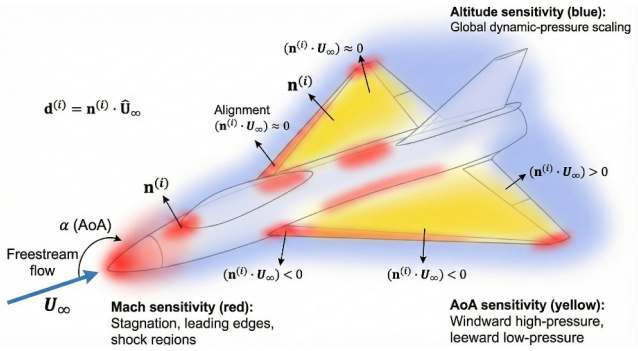


FIGURE 3 | Physical intuition of the CSLM module. The freestream direction and angle of attack define the local flow incidence relative to the surface normals. The red, yellow and blue regions schematically indicate areas mainly sensitive to Mach number, angle of attack and altitude, respectively.

Third, the FiLM transformation $\tilde{h}^{(i)} = \gamma^{(i)} \odot \phi(\hat{x}^{(i)}) + \beta^{(i)}$ uses these sensitivities to modulate the geometric features in a physically interpretable way: the scale term $\gamma^{(i)}$ acts as a learnt multiplicative factor that can capture compressibility- and dynamic-pressure-related amplification or attenuation of the local response, whereas the shift term $\beta^{(i)}$ introduces additive offsets that are well suited to represent baseline pressure shifts and redistribution effects associated with changes in AoA and boundary-layer state. As a result, the CSLM output $\left\{\tilde{h}^{(i)}\right\}$ encodes not only the local shape but also how that shape is expected to respond under the current flow conditions. These condition-modulated features are then propagated by the graph attention and multiscale transformer modules to enforce local consistency and global aerodynamic balance across the entire configuration.

3.3 | Flow-Aligned Graph Attention Layer

After CSLM, each surface point is represented by a FiLM-modulated feature vector $\left\{\tilde{h}^{(i)}\right\}_{i=1}^{N}$, which already encodes its local geometry and condition-sensitive response to Mach number, angle of attack and altitude. However, the surface pressure field is not determined pointwise: due to the elliptic nature of the pressure equation and the combined action of viscous and inviscid transport, the pressure at a given location is strongly coupled to its neighbours along the aircraft surface. To represent this spatial coupling in a physically meaningful way, APTNet employs a FAGAT operating on a surface adjacency graph.

3.3.1 | *Flow-Aligned Anisotropic Graph Construction*

Considering that supersonic flows exhibit pronounced anisotropy—flow quantities tend to vary smoothly along streamlines but can undergo sharp discontinuities when crossing shock waves—an isotropic k-NN graph built under the Euclidean metric may introduce spurious cross-shock connections and cause nonphysical smoothing. To alleviate this, we construct a flow-aligned anisotropic surface graph using a weighted flow distance. Let $\mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^3$ denote the coordinates of two surface points and $\mathbf{r}_{ij} = \mathbf{x}_j - \mathbf{x}_i$ be their relative displacement. Using the global freestream velocity vector $\mathbf{U}_\infty$, we define the normalised freestream direction

$$\hat{\mathbf{u}}_\infty = \frac{\mathbf{U}_\infty}{\|\mathbf{U}_\infty\|}. \tag{12}$$

We then decompose $\mathbf{r}_{ij}$ into components parallel and orthogonal to the freestream direction:

$$\mathbf{r}_{\parallel} = (\mathbf{r}_{ij} \cdot \hat{\mathbf{u}}_\infty)\hat{\mathbf{u}}_\infty, \qquad \mathbf{r}_{\perp} = \mathbf{r}_{ij} - \mathbf{r}_{\parallel}. \tag{13}$$

Based on this decomposition, we define a flow-aligned weighted distance

$$D_{\text{flow}}(i,j) = \sqrt{\omega_{\parallel}\|\mathbf{r}_{\parallel}\|^2 + \omega_{\perp}\|\mathbf{r}_{\perp}\|^2}, \tag{14}$$

where $\omega_{\parallel}$ and $\omega_{\perp}$ are directional scaling factors with $\omega_{\perp} > \omega_{\parallel}$. This choice penalises offsets orthogonal to the freestream more strongly than along-streamline offsets, thereby reshaping the local receptive field from a sphere into an ellipsoid elongated along $\hat{\mathbf{u}}_{\infty}$. The neighbour set N_i of node i is obtained by selecting the k points with the smallest $D_{\text{flow}}(i,j)$. In practice, we use the union of directed k-NN edges to form an undirected graph. This flow-aligned construction encourages information propagation preferentially along streamlines whilst suppressing spurious cross-shock feature mixing, thereby improving the preservation of strong-gradient structures such as shock footprints. As illustrated in Figure 4, an isotropic Euclidean kNN neighbourhood can easily connect points across a shock interface, whereas the proposed flow-aligned construction reshapes the receptive field into an ellipsoid elongated along $\hat{\mathbf{u}}_{\infty}$, thereby discouraging cross-shock feature aggregation.

3.3.2 | *Graph Attention Message Passing*

On this graph, the FAGAT updates each node by aggregating information from its neighbours through attention-weighted message passing. Let N_i denote the neighbour set of node i. For the k-th attention head, we first apply a learnable linear projection to obtain $h_i^{(k)}$ from $\tilde{h}^{(i)}$ and then compute the updated feature as

$$h_i^{(k)\prime} = \sum_{j \in N_i} \alpha_{ij}^{(k)} h_j^{(k)}, \tag{15}$$

where $\alpha_{ij}^{(k)}$ are normalised attention coefficients learnt from the feature pairs $\left(h_i^{(k)}, h_j^{(k)}\right)$. These coefficients can be interpreted as the strength of aerodynamic influence from location j to location i: larger values naturally emerge along directions where pressure changes are strongly coupled, such as across shock footprints, along highly curved surfaces or within attached boundary-layer regions, whereas weaker couplings are assigned to aerodynamically less relevant neighbours.

The outputs of all attention heads are concatenated and linearly combined to form the updated node feature. In this way, the FAGAT module enforces local pressure continuity and smooth propagation of condition-modulated information along the surface topology. It acts as a topology-aware physics-inspired preprocessor for the subsequent multiscale transformer block, which then captures longer-range aerodynamic coupling across the entire configuration, including global lift distribution, shock-system interaction and pressure-centre movement.

3.4 | Multi-Scale Transformer Block

Although the graph attention layer enforces local topology-aware propagation of condition-modulated features, many aerodynamic phenomena are governed by long-range interactions that extend well beyond immediate surface neighbourhoods. Examples include wing–body interference, shock–shock and shock–boundary-layer interactions as well as global lift redistribution and pressure-centre shifts. To capture such nonlocal coupling, APTNet employs a multiscale transformer block acting on the set of node features output by the FAGAT module.

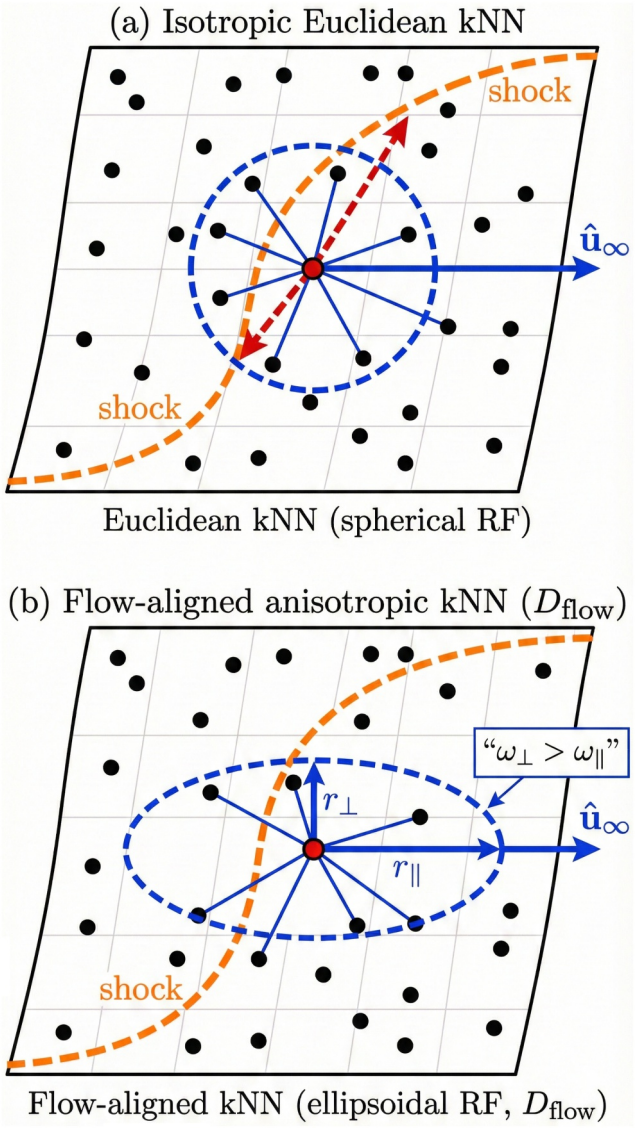


FIGURE 4 | Comparison of neighbourhood construction strategies for surface-graph learning in supersonic flows. (a) Isotropic Euclidean kNN produces a spherical receptive field. (b) The proposed flow-aligned anisotropic kNN reshapes the receptive field into an ellipsoid elongated along the freestream direction, thereby reducing cross-shock connections.

Let $\{z^{(i)}\}_{i=1}^{N}$ denote the updated node features after FAGAT, where each $z^{(i)} \in \mathbb{R}^H$ encodes the locally propagated condition-sensitive state at surface point i. We treat these features as a set of tokens and form a sequence $\boldsymbol{Z} = \left[\boldsymbol{z}^{(1)}, \boldsymbol{z}^{(2)}, ..., \boldsymbol{z}^{(N)}\right] \in \mathbb{R}^{N \times H}$. A standard multihead self-attention layer is then applied:

$$\text{Attn}(\boldsymbol{Z}) = \text{softmax}\left(\frac{\boldsymbol{Q}\boldsymbol{K}^{\top}}{\sqrt{H}}\right)\boldsymbol{V}, \tag{16}$$

where $\boldsymbol{Q} = \boldsymbol{Z}\boldsymbol{W}_Q, \boldsymbol{K} = \boldsymbol{Z}\boldsymbol{W}_K$ and $\boldsymbol{V} = \boldsymbol{Z}\boldsymbol{W}_V$ are the query, key and value projections and $\boldsymbol{W}_Q, \boldsymbol{W}_K, \boldsymbol{W}_V$ are learnable matrices. Multihead attention is used in practice, and the attention weights quantify how strongly the pressure state at one surface location is influenced by that at other locations across the configuration.

To enhance the ability of the network to capture interactions at different spatial scales, we stack multiple transformer encoder layers and interpret the resulting hierarchy as a multiscale representation. In the lower layers, self-attention mainly couples nearby points with similar geometric and condition-modulated features, reinforcing regional patterns such as local shock structures or high-curvature loading. As depth increases, the effective receptive field expands, enabling the model to capture global interactions such as wing–body coupling, downstream wake effects and overall lift and moment redistribution. Residual connections and position-wise feed-forward networks are employed in each layer, following standard transformer practice, to stabilise training and increase expressiveness.

Physically, the multiscale transformer block provides a complementary role to the FAGAT module. Although FAGAT focuses on enforcing local consistency along the surface mesh and modelling short-range pressure propagation constrained by topology, the transformer captures long-range configuration-wide couplings that arise from the elliptic nature of the pressure field and the global structure of the flow. The final transformer outputs $\{\hat{z}^{(i)}\}_{i=1}^{N}$ serve as globally informed, condition-aware representations of the surface points, which are then fed into a lightweight prediction head to regress the surface pressure at each location.

3.5 | Implementation Details

For reproducibility, we further specify the architectural settings and training hyperparameters of APTNet. Unless otherwise stated, the hidden feature dimension is set to $H = 192$ throughout the network. The detailed architectural configurations of the main modules are summarised in Table 1, and the training hyperparameters are listed in Table 2.

TABLE 1 | Architectural hyperparameters of APTNet.

Module/Item	Configuration
Hidden dimension	$H = 192$ throughout the network
Point encoder	Two-layer MLP: $3 \to 96 \to 192$; each layer followed by LayerNorm and ReLU; dropout rate $= 0.1$ after the first layer
Condition encoder	Each scalar condition (Mach number, AoA and altitude) is independently projected from $1 \to 64$
Local sensitivity predictor	Three-layer MLP: $385 \to 192 \to 96 \to 3$; LayerNorm and ReLU between hidden layers; dropout rate $= 0.1$ and sigmoid activation on output
FiLM modulation MLP	Two-layer MLP: $384 \to 192 \to 384$; the output is split into scale and shift vectors
Graph construction	Flow-aligned anisotropic kNN graph with $k = 16$ neighbours
FAGAT	4 attention heads, 48 dimensions per head; concatenated output projected back to 192 dimensions and residual connection used
Directional weights	$\omega_{\parallel} = 1.0, \omega_{\perp} = 4.0$
Multi-scale transformer	4 encoder layers; 8 attention heads; hidden dimension $= 192$; feed-forward dimension $= 384$ and dropout rate $= 0.1$
Prediction head	Three-layer MLP: $192 \to 192 \to 96 \to 1$; LayerNorm, ReLU and dropout rate $= 0.1$ between adjacent linear layers

4 | Experiments and Analysis

4.1 | The Dataset and the Experimental Environment

To evaluate the performance of the proposed model, we constructed a dataset based on four representative aircraft geometries, namely blunt cone, double ellipsoid, double blunt cone and aerospacecraft. The flight condition parameters are summarised in Table 3. We used CFD simulations to compute the surface-pressure distributions of these aircraft under different operating conditions.

CFD can simulate various complex geometric shapes and flow conditions, generating a wide range of datasets with high accuracy and covering extensive operating scenarios. This provides diverse reference data for training and evaluation and supports the assessment of the model's generalisation capability. Based on these advantages, we chose CFD to generate the ground-truth data. In total, 880 samples were obtained for model development and evaluation. To further clarify the experimental protocol, the dataset split was performed in a geometry-aware manner. Specifically, the aerospacecraft configuration was excluded from

TABLE 2 | Training hyperparameters of APTNet.

Item	Setting
Optimiser	AdamW
Initial learning rate	1×10^{-3}
Weight decay	1×10^{-4}
Batch size	4
Training epochs	300
Learning-rate schedule	Cosine annealing
Warmup	10 epochs
Early stopping	Based on validation MRE

TABLE 3 | Flight conditions and their values.

Altitude (km)	Mach number (Ma)	Angle of attack (deg)
7.0	3.0	− 20
7.5	3.5	− 15
8.0	4.0	− 10
8.5	4.5	− 5
9.0	5.0	0
⋯	⋯	5
⋯	⋯	10
⋯	⋯	15
20.0	20.0	20

the training stage and reserved as a held-out test geometry to evaluate cross-geometry generalisation. The remaining three geometries were used to construct the training and validation subsets. Overall, the full dataset was partitioned into training, validation and test subsets with an approximate ratio of 7:1:2. Therefore, the test set evaluates not only predictive performance under varying operating conditions but also the robustness of APTNet on an unseen aircraft geometry.

The three-dimensional compressible Navier–Stokes equations in Cartesian coordinates can be written as follows:

$$\frac{\partial \boldsymbol{Q}}{\partial t}+\frac{\partial \boldsymbol{E}}{\partial x}+\frac{\partial \boldsymbol{F}}{\partial y}+\frac{\partial \boldsymbol{G}}{\partial z}=\frac{1}{\mathrm{Re}}\left(\frac{\partial \boldsymbol{E}_v}{\partial x}+\frac{\partial \boldsymbol{F}_v}{\partial y}+\frac{\partial \boldsymbol{G}_v}{\partial z}\right) \tag{17}$$

In Equation (17), Re is the Reynolds number, which represents the ratio of inertial to viscous forces. The Cartesian coordinate system defines the spatial coordinates (x, y, z) and t denotes time. The conserved state vector is denoted by Q, whereas the convective flux vectors are denoted by E, F and G. The corresponding diffusive flux vectors are E_v, F_v and G_v, respectively and are given by

$$\boldsymbol{Q}=\begin{bmatrix}\rho\\ \rho u\\ \rho v\\ \rho w\\ e\end{bmatrix},\quad \boldsymbol{E}=\begin{bmatrix}\rho u\\ \rho u^2+p\\ \rho uv\\ \rho uw\\ (e+p)u\end{bmatrix},\quad \boldsymbol{E}_{\mathrm{v}}=\begin{bmatrix}0\\ \tau_{xx}\\ \tau_{xy}\\ \tau_{xz}\\ \tau_{xx}u+\tau_{xy}v+\tau_{xz}w+q_x\end{bmatrix},$$

$$\boldsymbol{F}=\begin{bmatrix}\rho v\\ \rho uv\\ \rho v^2+p\\ \rho vw\\ (e+p)v\end{bmatrix},\quad \boldsymbol{F}_{\mathrm{v}}=\begin{bmatrix}0\\ \tau_{xy}\\ \tau_{yy}\\ \tau_{yz}\\ \tau_{xy}u+\tau_{yy}v+\tau_{yz}w+q_y\end{bmatrix},$$

$$\boldsymbol{G}=\begin{bmatrix}\rho w\\ \rho uw\\ \rho vw\\ \rho w^2+p\\ (e+p)w\end{bmatrix},\quad \boldsymbol{G}_{\mathrm{v}}=\begin{bmatrix}0\\ \tau_{xz}\\ \tau_{yz}\\ \tau_{zz}\\ \tau_{xz}u+\tau_{yz}v+\tau_{zz}w+q_z\end{bmatrix}. \tag{18}$$

In Equation (18), the Cartesian velocity components u, v and w represent the velocity in the x, y and z directions, respectively. Similarly, the variables ρ, p and e denote the fluid density, pressure and total energy per unit volume, respectively. Finally, the stress in the j direction acting on a plane perpendicular to the i axis is represented by τ_{ij}. The heat-flux components in the x, y and z directions are denoted by q_x, q_y and q_z, respectively.

For CFD data generation, all simulations were performed in ANSYS fluent 2023 R1 using the three-dimensional density-based implicit solver for compressible flow. The governing equations were solved in the RANS framework rather than an inviscid Euler setting, and turbulence was modelled using the k–ω SST model. The energy equation was enabled, and air was treated as an ideal gas. The computational domain was discretised using an unstructured mixed-cell mesh. In terms of mesh strategy, local refinement was applied near the aircraft surface and in geometrically complex or high-gradient regions to better resolve the surface-pressure distribution. The far-field conditions were imposed through pressure-far-field boundaries according to the specified flight conditions, and symmetry boundaries were used where appropriate. For the numerical solution, gradients were evaluated using the least-squares cell-based method, whereas the flow, turbulent kinetic energy and specific dissipation rate equations were discretised using second-order upwind schemes. AUSM flux splitting was adopted in the density-based formulation. Convergence was judged from the residual history, with a target residual level of 10^{-3} for all solved equations, together with stabilisation of the surface-pressure distribution.

The CFD reference data were generated in ANSYS Fluent, whereas model training and inference were conducted on a workstation equipped with three NVIDIA RTX A6000 GPUs.

4.2 | Model Performance Evaluation

Given that surface-pressure values span a wide dynamic range across geometries and operating conditions, we adopt the mean relative error (MRE) as both the training loss and the primary evaluation metric. The MRE is defined as

$$\mathrm{MRE}=\frac{1}{N}\sum_{i=1}^{N}\left|\frac{\hat{y}_i-y_i}{y_i}\right|, \tag{19}$$

where N is the number of surface elements, y_i denotes the CFD reference pressure and $\hat{y}_i$ is the model prediction. As a dimensionless measure, MRE provides an intuitive percentage-type error that remains comparable across cases with different pressure magnitudes.

Table 4 reports quantitative results on the four representative geometries in Figures 5 and 6. Across all configurations, APTNet achieves the lowest MRE among the compared learning-based baselines, including the point-based SA-HFNet and the graph-based RegDGCNN. This indicates that, for surface pressure fields, improving geometric encoding alone (pure point networks) or introducing topology via graph message passing (graph networks) is still insufficient if the flow-condition influence is injected in a largely global or uniform manner. In contrast, APTNet's combination of condition-sensitive local

TABLE 4 | Comparing the performance of our model with other models.

Method	MRE (%) on different geometries				Mean inference time	Speedup over CFD
	Blunt cone	Double blunt cone	Double ellipsoid	Aerospacecraft		
PointNet++ [15]	5.78%	6.28%	6.32%	6.91%	87.2 ms	≈6.8 × 10^3
SA-HFNet [34]	5.61%	5.21%	6.72%	4.97%	**64.6 ms**	≈**9.2** × $\mathbf{10^3}$
PointConvAttn [17]	3.93%	4.67%	5.21%	4.62%	75.5 ms	≈7.9 × 10^3
PointNeXtReg+ [19]	4.06%	4.59%	5.62%	4.95%	83.3 ms	≈7.1 × 10^3
RegDGCNN [18]	5.22%	4.98%	4.88%	5.37%	105.6 ms	≈5.6 × 10^3
APTNet (ours)	**3.68%**	**4.57%**	**4.51%**	**4.33%**	93.1 ms	≈6.4 × 10^3

Note: Bold values indicate the most favorable results in each column.

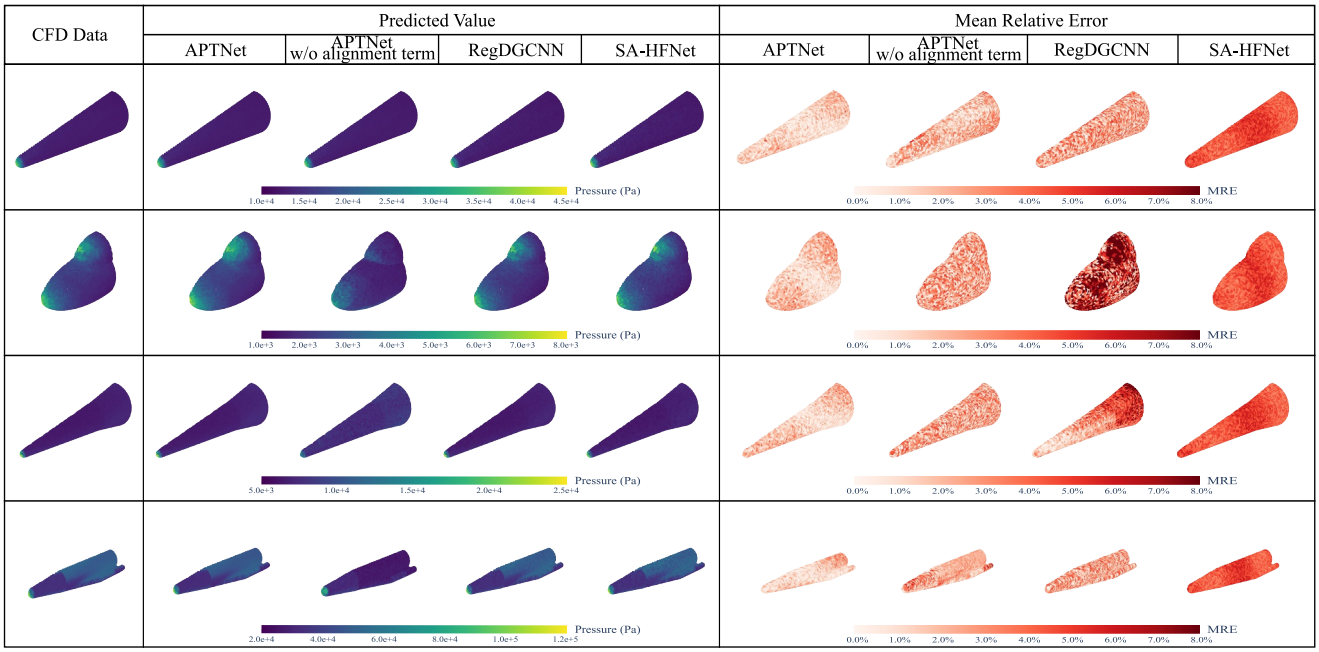


FIGURE 5 | Qualitative comparison of CFD references and model predictions on representative test cases. For each geometry, we show the CFD pressure field, predictions from RegDGCNN, SA-HFNet, APTNet and the ablation variant without the alignment term together with the corresponding point-wise relative error maps.

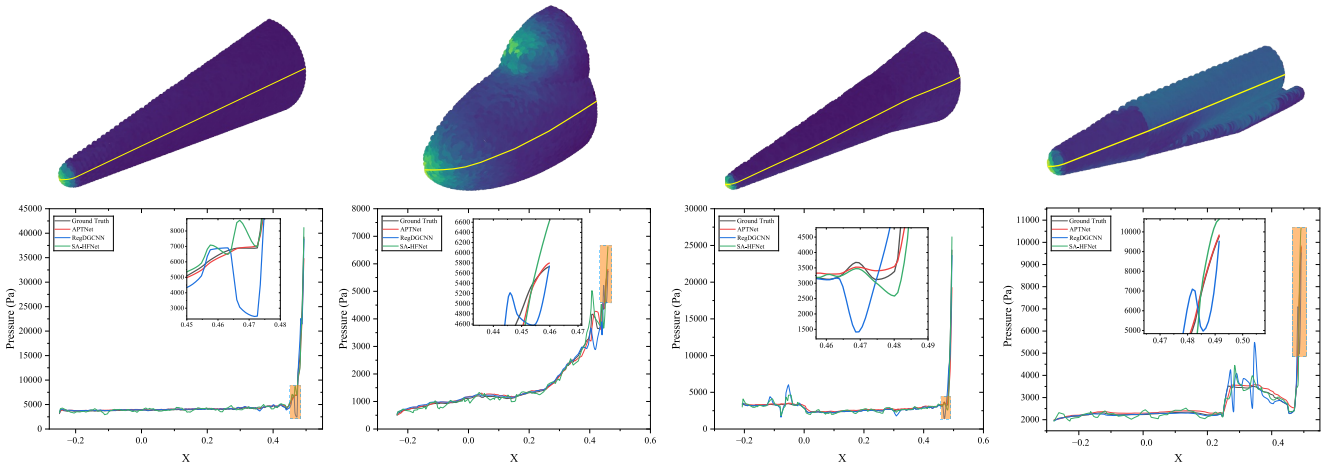


FIGURE 6 | The pressure profiles along the surface slices highlighted by the yellow lines comparing CFD with APTNet, RegDGCNN and SA-HFNet.

modulation with topology-aware propagation yields more accurate and more transferable pressure predictions across both geometry variants and operating conditions. Meanwhile, the mean inference time of APTNet remains on the order of 10^2 ms per case, which is comparable to other deep-learning surrogates. Under the present Fluent setting, the average computational time of a standard CFD simulation is approximately 10 min per case. As summarised in Table 4, APTNet achieves a speedup of about 6.4×10^3 relative to CFD whilst maintaining superior prediction accuracy across the four geometries. Such an efficiency advantage makes the proposed surrogate particularly suitable for iterative evaluation and rapid aerodynamic analysis.

To highlight method differences under challenging conditions, for each geometry we select the test sample with the largest MRE and visualise the CFD reference, model predictions and point-wise relative errors in Figure 5. Compared with RegDGCNN and SA-HFNet, APTNet better preserves sharp pressure transitions and localised high-gradient regions, whereas the baselines tend to exhibit larger contiguous error patches and over-smoothed structures in these areas. Figure 5 also includes an ablation variant without the flow–geometry alignment term. Removing this alignment cue leads to visibly degraded localisation of the error patterns, suggesting that the alignment term provides an effective geometric prior for organising spatially heterogeneous condition responses on complex surfaces.

In addition to the full-field visualisations, the slice-wise profiles in Figure 6 further confirm that APTNet tracks peak locations and amplitudes more faithfully than the baselines, particularly in regions where the pressure varies rapidly along the surface. To further quantify the ability of different models to preserve sharp pressure transitions near shock-related regions, we introduce a shock-region mean relative error (SR-MRE) metric based on the slice-wise pressure profiles in Figure 6. For each selected surface slice, let p_i denote the CFD reference pressure at the i-th sampled point along the slice and let $\hat{p}_i$ denote the corresponding model prediction. We first compute the reference pressure-gradient magnitude along the slice as

$$g_i = \left| \frac{p_{i+1} - p_i}{s_{i+1} - s_i} \right|, \tag{20}$$

where s_i is the curvilinear coordinate of the i-th point on the slice. Based on the CFD reference, the shock-related region Ω_s is defined as the set of points whose gradient magnitude ranks in the top 10% along that slice, that is,

$$\Omega_s = \left\{ i \mid g_i \text{ belongs to the top } 10\% \text{ of } \left\{ g_j \right\} \right\}. \tag{21}$$

The corresponding shock-region mean relative error is then defined as

$$\text{SR-MRE} = \frac{1}{|\Omega_s|} \sum_{i \in \Omega_s} \left| \frac{\hat{p}_i - p_i}{p_i} \right|. \tag{22}$$

A lower SR-MRE indicates better preservation of sharp pressure variations and less nonphysical smoothing near shock-related structures. This metric is particularly relevant to FAGAT, whose anisotropic flow-aligned graph construction is designed to suppress spurious cross-shock feature mixing.

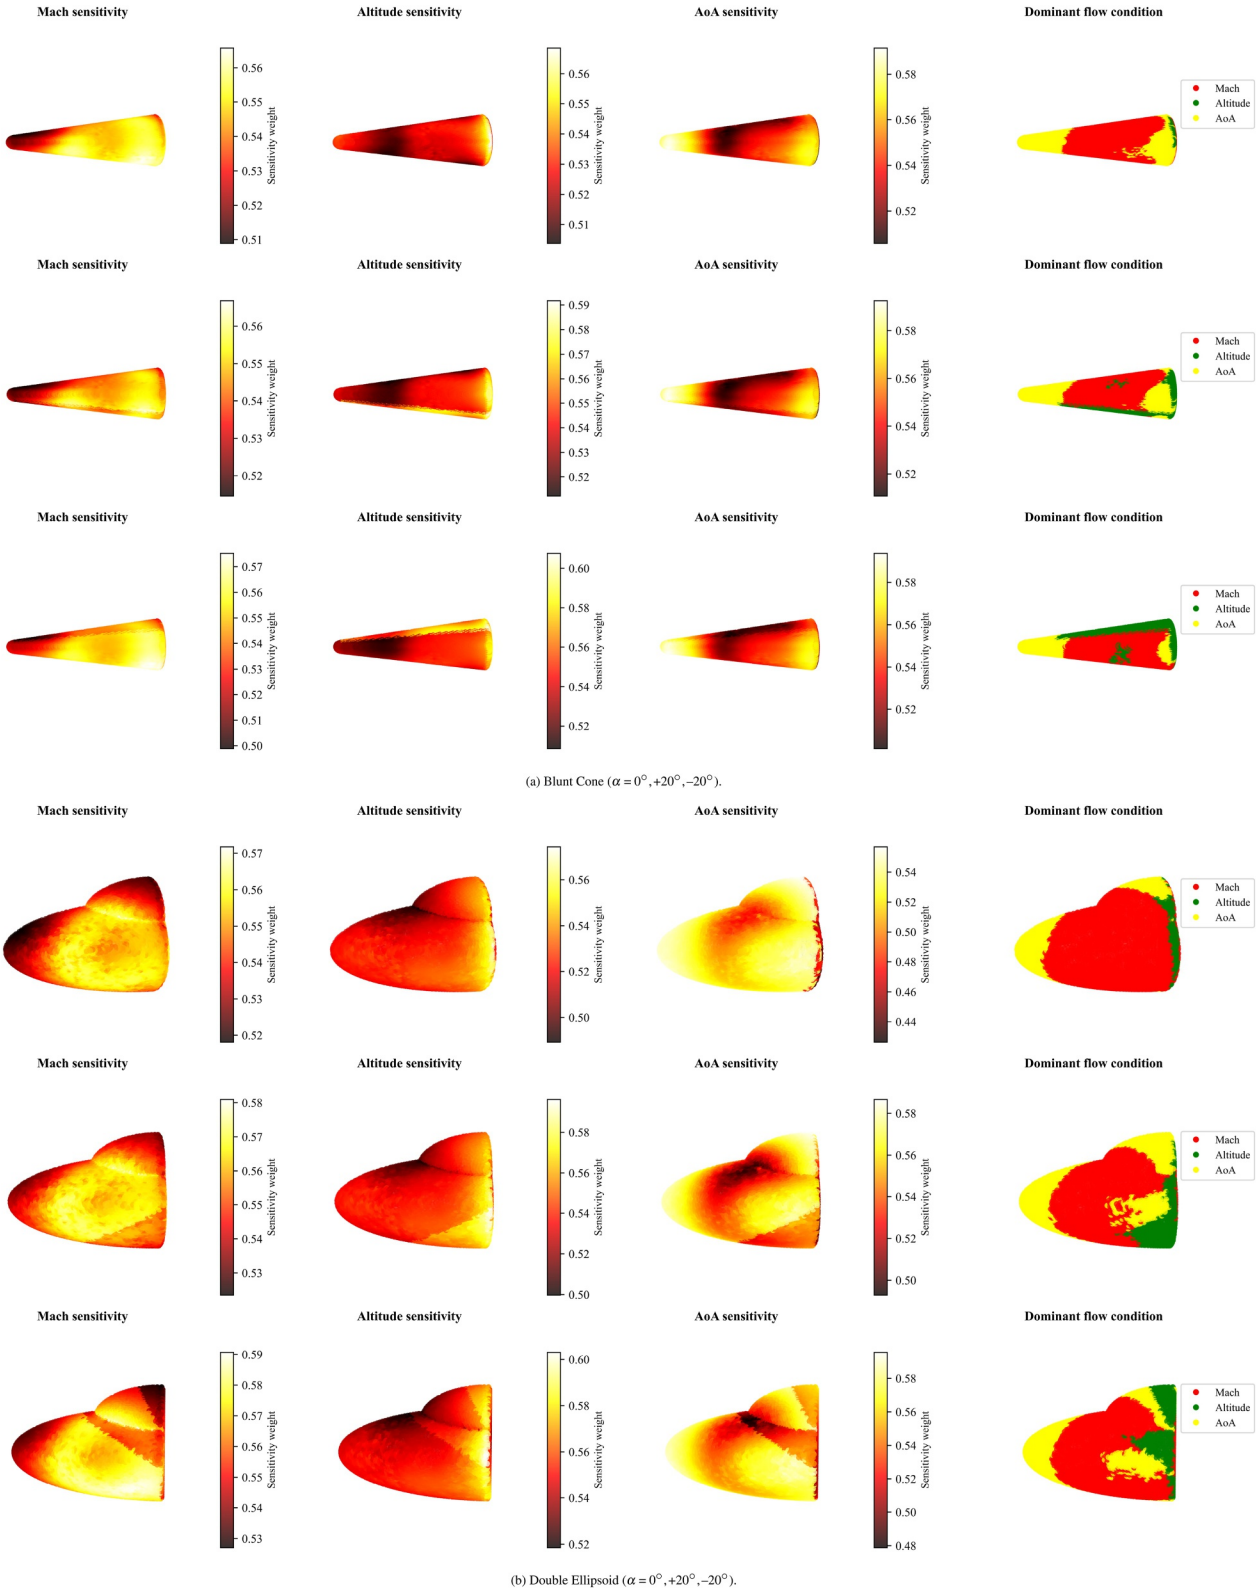
As shown in Table 5, APTNet consistently achieves the lowest SR-MRE across all four geometries. Compared with the graph-based baseline RegDGCNN, the proposed method yields substantially lower errors in the high-gradient shock-related regions indicating a stronger ability to preserve sharp pressure transitions. Moreover, the full APTNet also outperforms the ablation variant without FAGAT, which quantitatively confirms the effectiveness of the flow-aligned graph attention layer in suppressing nonphysical cross-shock smoothing and improving the reconstruction of shock-related pressure variations. These results provide quantitative support for the visual comparisons in Figure 6.

To examine the physical plausibility of the learnt CSLM, we visualise the condition-sensitivity fields on four representative aircraft geometries under three angles of attack, $\alpha \in \{0°, +20°, -20°\}$ (Figure 7–9). For each case, the first three panels show the normalised point-wise sensitivity weights $w_{\text{Ma}}^{(i)}$, $w_{Alt}^{(i)}$ and $w_{\text{AoA}}^{(i)}$ over the surface, whereas the last panel highlights the dominant controlling condition at each point. Several consistent, physically interpretable patterns emerge across geometries and AoAs. First, the Mach sensitivity tends to concentrate around stagnation zones and leading-edge regions, where compressibility and shock-related effects are expected to be most pronounced. Second, the AoA sensitivity exhibits a clear windward–leeward asymmetry, reflecting the dominant role of incidence in redistributing surface loading and generating lift-related pressure contrasts; moreover, as $|\alpha|$ increases (e.g., $\pm 20°$), the sensitivity maps become markedly more spatially asymmetric, and reversing the sign of α leads to a corresponding windward–leeward role exchange rather than a fixed spatial partition. Third, the altitude sensitivity fields are generally smoother and more spatially distributed, which is consistent with their role in modulating the overall dynamic-pressure level instead of introducing highly localised effects. At $\alpha = 0°$, the dominant-condition maps are comparatively more balanced and less segregated in line with reduced incidence-driven asymmetry. For geometries with severe self-occlusion or concave structures (e.g., the aerospacecraft case), windward and leeward surfaces are visualised separately; nevertheless, the cross-AoA consistency remains evident in that the windward-side patterns at $+20°$ qualitatively correspond to the leeward-side patterns at $-20°$ (and vice versa), whilst retaining geometry-specific localised

TABLE 5 | Shock-region mean relative error (SR-MRE, %) on the selected slices in Figure 6.

Method	Blunt cone	Double blunt cone	Double ellipsoid	Aerospacecraft
RegDGCNN	5.96	5.48	5.31	5.72
APTNet w/o FAGAT	4.98	4.73	4.88	4.65
APTNet (ours)	**4.27**	**4.11**	**4.36**	**4.08**

Note: Lower is better. Bold values indicate the most favorable results in each column.



(a) Blunt Cone ($\alpha = 0°, +20°, -20°$).

(b) Double Ellipsoid ($\alpha = 0°, +20°, -20°$).

FIGURE 7 | CSLM condition-sensitivity visualisations for blunt cone and double ellipsoid under three angles of attack.

structures near junctions, edges or high-curvature regions. Overall, the observed AoA-driven asymmetry, sign-reversal consistency and geometry-adaptive spatial organisation collectively support that CSLM captures physically meaningful, spatially heterogeneous condition dependence across both operating conditions and aircraft geometries.

The dominant-condition maps further indicate that different surface regions of the same configuration are governed by different flow parameters (e.g., nose and leading-edge areas are largely Mach-dominated, whereas lifting surfaces are often AoA-dominated). These qualitative patterns align with aerodynamic intuition and support the design goal of CSLM:

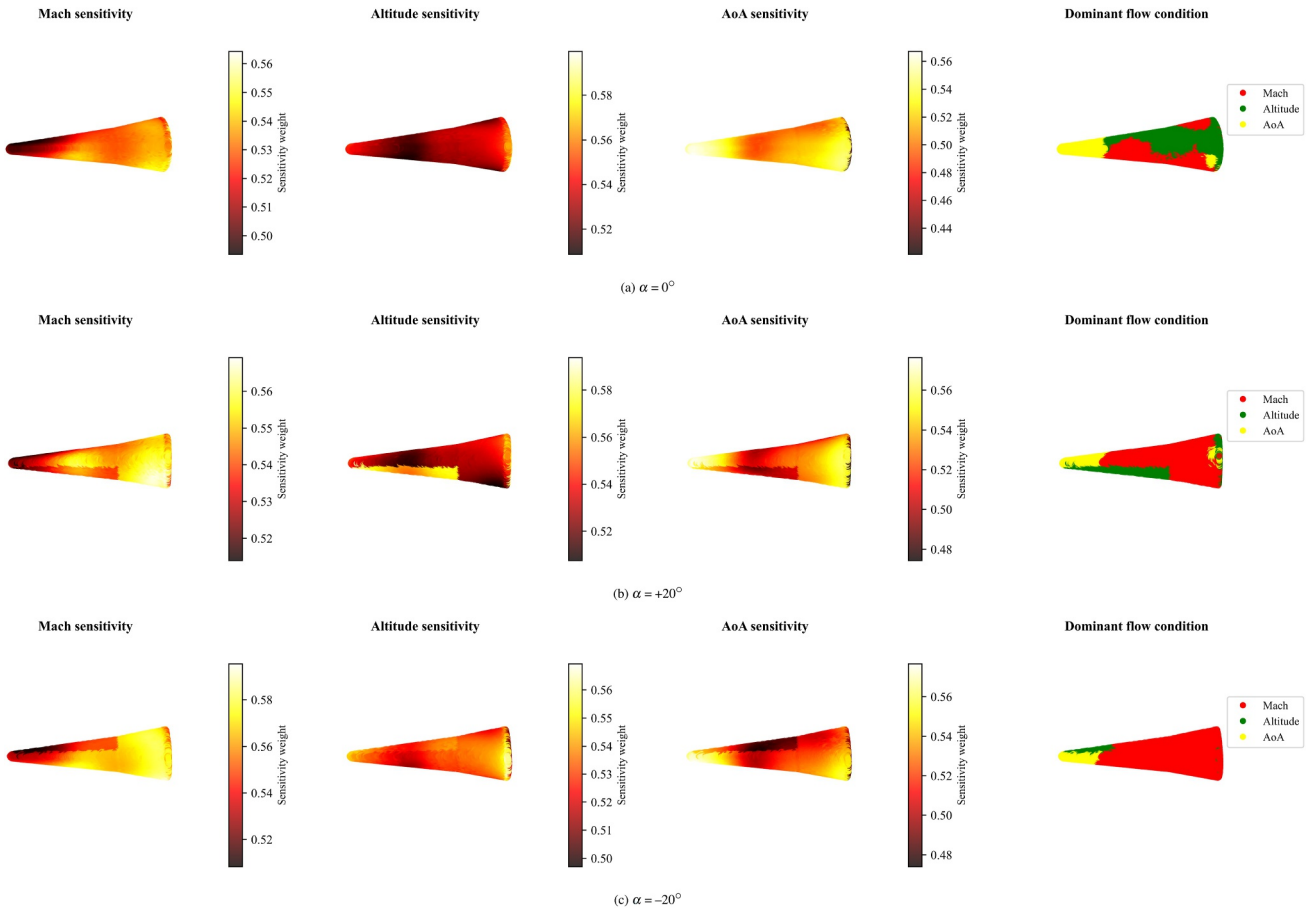


FIGURE 8 | CSLM condition-sensitivity visualisations for double blunt cone under three angles of attack.

learning spatially heterogeneous, physically plausible condition dependence rather than assuming a uniform response to changes in flow conditions.

4.3 | Ablation Experiment

To evaluate the contribution of each key component in APTNet, we conduct a systematic ablation study. In particular, we focus on the CSLM module and the FAGAT layer, which are responsible for condition-dependent local modelling and topology-aware feature propagation, respectively. We construct the following model variants:

1. Full model (APTNet). The complete architecture with both CSLM and FAGAT enabled.
2. No CSLM (APTNet-CSLM). The CSLM module is removed and flow conditions are fused with geometric features via simple concatenation, resulting in a purely global fusion of shape and condition.
3. No FAGAT (APTNet-FAGAT). The graph attention layer is removed and the condition-modulated point features are fed directly into the multiscale transformer blocks, without explicit topology-aware local propagation.
4. No CSLM and FAGAT (APTNet–CSLM–FAGAT). Both CSLM and FAGAT are removed, leaving a baseline architecture that relies solely on a transformer operating on position-encoded point features with global condition concatenation.

All variants are trained on the same dataset with identical optimisation settings, learning rate schedules, and regularisation strategies to ensure a fair comparison. The quantitative results on the test set are summarised in Table 6, where we report the MRE and parameter counts for each model.

From these experiments, several observations can be made.

4.3.1 | *Impact of Removing CSLM*

When CSLM is removed (APTNet–CSLM), the MRE increases from 3.68% to 4.57% corresponding to a relative degradation of about 24%. In this variant, flow conditions are only injected through a global concatenation, so all surface points share the same condition representation. As a result, the model loses the ability to represent the spatially nonuniform sensitivity of the pressure field to Mach number, AoA and altitude. In particular, stagnation regions, windward and leeward surfaces and shock-sensitive areas are no longer allowed to respond differently to the same condition change, which leads to larger errors in regions with strong local aerodynamic variation.

4.3.2 | *Impact of Removing FAGAT*

When the FAGAT layer is removed (APTNet–FAGAT), the MRE increases to 4.51%, a relative degradation of about 22%. In this case, the condition-modulated features produced by CSLM are no longer propagated through an explicit adjacency graph, and the model must rely solely on the transformer to capture

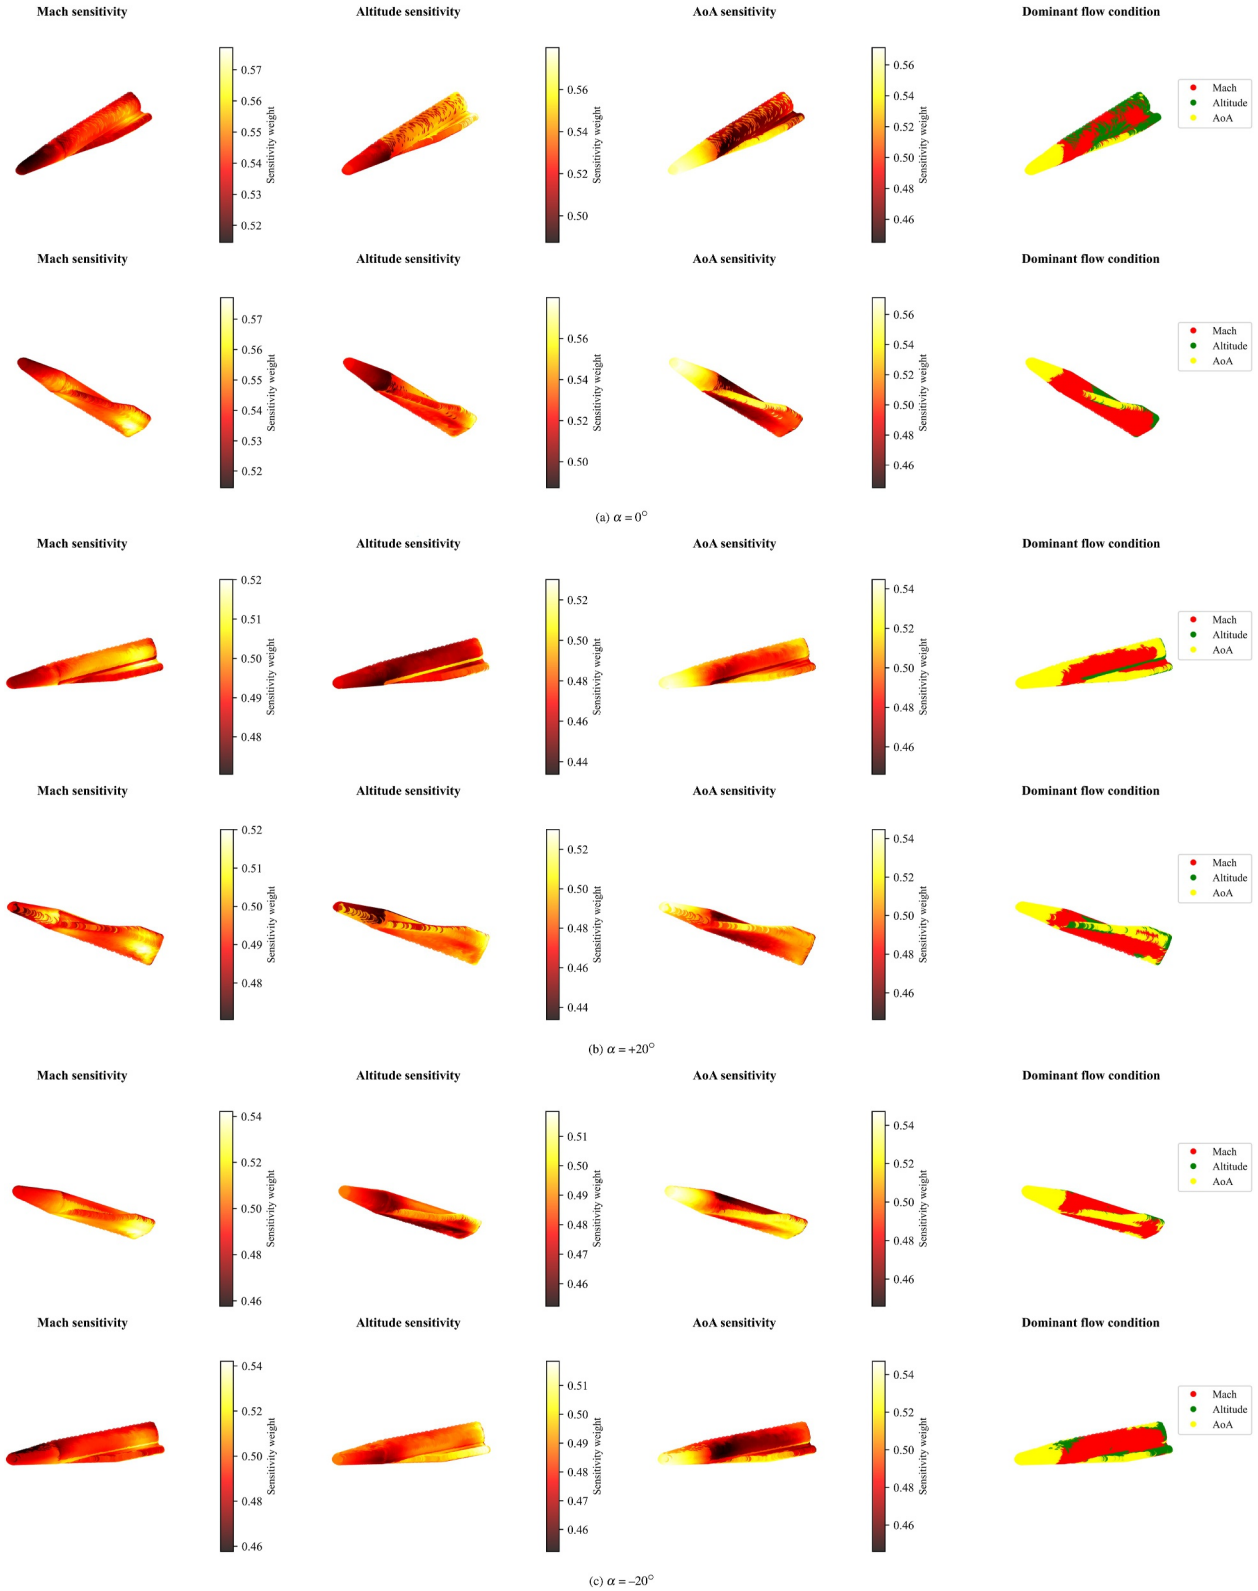


FIGURE 9 | CSLM condition-sensitivity visualisations for aerospacecraft. Visualisations of the aerospacecraft geometry under two complementary viewpoints to provide more complete coverage of the surface features.

both local and global interactions. The performance drop indicates that the FAGAT layer plays an important role in enforcing local pressure consistency and exploiting the mesh topology. In particular, without FAGAT the model is less capable of handling complex geometric structures, such as leading edges, wing–tail junctions and regions with large curvature gradients, where local neighbourhood interactions are crucial.

4.3.3 | *Synergy Between CSLM and FAGAT*

When both CSLM and FAGAT are removed simultaneously (APTNet–CSLM–FAGAT), the degradation is the most severe: the MRE rises to 4.95%, which is about 35% worse than the full model. This indicates that CSLM and FAGAT are not simple interchangeable add-ons, but rather complementary components. CSLM provides a physically motivated per-point condition-sensitivity field, whereas FAGAT propagates this information along the surface topology, shaping how local responses interact and equilibrate. The strong degradation of the baseline variant confirms that accurate surface pressure prediction requires both condition-sensitive local modelling and topology-aware neighbourhood propagation, providing empirical support for the design of APTNet.

4.4 | Comparison With Other Methods

To further assess the effectiveness of APTNet, we compare it with several representative deep-learning surrogates for three-dimensional aerodynamics, including a PointNet++-based pressure predictor [15], the shape-aware SA-HFNet [34], the point-convolution-attention architecture PointConvAttn [17], the PointNeXtReg + transfer-learning framework [19] and the RegDGCNN graph network originally developed in the DrivAerNet work for automotive aerodynamics [18]. Table 4 reports the MRE on the four aircraft configurations as well as the mean inference time.

Overall, APTNet achieves the lowest MRE on all shapes. On the blunt cone and aerospacecraft configurations, APTNet reduces the error by approximately 6%–8% relative to the strongest competing method, whereas on the double blunt-cone and double ellipsoid it reaches a comparable accuracy to PointNeXtReg+. These gains are most pronounced for the aerospacecraft, where the surface geometry exhibits strong curvature, windward–leeward asymmetry and complex shock-interaction patterns. In such cases, pure point-based networks (PointNet++ and SA-HFNet) suffer from larger errors because they treat flow conditions as globally shared inputs and do not explicitly encode mesh topology. PointConvAttn and PointNeXtReg + alleviate this limitation by incorporating local neighbourhoods or stronger geometric encoders, but they still rely on global shape-condition fusion. RegDGCNN, as a graph-based model, better exploits surface connectivity; however, it also injects flow conditions only globally and was originally designed for integrated coefficients rather than pointwise pressure fields.

By contrast, APTNet combines three complementary mechanisms: the CSLM module builds a per-point, physically motivated condition-sensitivity field, the FAGAT layer propagates these modulated features along the surface topology and the multiscale transformer captures configuration-wide aerodynamic coupling. This design enables different regions (e.g., stagnation zones, lifting surfaces and leeward panels) to respond differently to Mach number, angle of attack and altitude, leading to systematically lower errors across all geometries. The qualitative comparisons in Figures 5 and 7–9 further confirm that APTNet better preserves shock locations and windward–leeward pressure contrasts than the baselines.

TABLE 6 | Ablation experiment.

Model	MRE (%)	Parameters (M)
APTNet	**3.68**	1.84
APTNet–CSLM	4.57	1.61
APTNet–FAGAT	4.51	1.54
APTNet–CSLM–FAGAT	4.95	**1.12**

Note: Bold values indicate the most favorable results in each column.

In terms of efficiency, APTNet is not the fastest model (93.1 ms mean inference time vs. 64.6 ms for SA-HFNet), but its runtime remains within the same order of magnitude as other deep-learning surrogates and is dominated by matrix-multiplication operations that are highly parallelisable on modern GPUs. Given the substantial accuracy improvements, especially on complex configurations, this moderate increase in inference time is acceptable for design- and analysis-oriented workflows where CFD simulations are several orders of magnitude slower.

5 | Conclusion

In this work, we proposed APTNet, a physics-informed learning framework for fast surface pressure prediction on high-speed vehicles under varying operating conditions. Motivated by the spatially heterogeneous dependence of surface pressure on Mach number, angle of attack and altitude, we revisited the limitation of global condition fusion and introduced a localised and interpretable mechanism for flow–geometry interaction.

To this end, APTNet incorporates a CSLM module that embeds flow–geometry interaction through a surface-normal–free-stream alignment indicator, $d^{(i)} = \mathbf{n}^{(i)} \cdot \hat{\mathbf{u}}_\infty$, and learns per-point sensitivity fields with respect to key conditions. These sensitivities are injected via FiLM-style modulation, enabling localised pressure amplification and redistribution while maintaining physical interpretability. In addition, a surface-based learning pipeline combining topology-aware local propagation with configuration-level coupling is designed to enhance accuracy

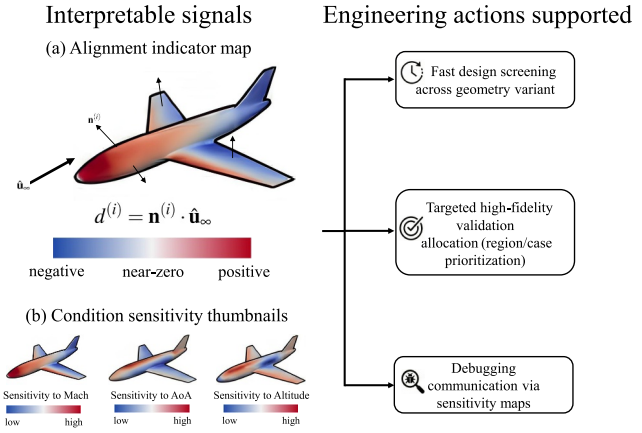


FIGURE 10 | Engineering value of the proposed interpretable signals. Left: (a) Visualisation of the flow–geometry alignment indicator on the surface (b) condition-sensitivity thumbnails. Right: Engineering actions supported.

and improve cross-geometry generalisation under varying operating conditions.

Experimental results show that APTNet consistently outperforms representative baselines in both accuracy and efficiency across all tested configurations and flow conditions. Beyond quantitative metrics, the model also provides interpretable intermediate signals that can be directly integrated into engineering workflows. As illustrated in Figure 10, these signals facilitate fast design screening across geometry variants, support targeted allocation of high-fidelity validation to representative regions and challenging cases and offer a practical interface for debugging and communicating model behaviour through condition-attributed patterns.

Future work will explore extending this framework to unsteady or manoeuvring flow conditions, incorporating additional physics-based constraints and coupling APTNet with gradient-based optimisation and digital design workflows to further accelerate aerodynamic analysis and design iteration.

Acknowledgements

This work is supported by the Research Fund of National Key Laboratory of Aerospace Physics in Fluids (Grant 2025-APFKFZD-30), the International Scientific and Technological Cooperation Projects of Sichuan Province (Grant 2025YFHZ0148), and the Central Government Guides Local Projects of China (Grant 2025ZYDF105).

Funding

Research Fund of National Key Laboratory of Aerospace Physics in Fluids, (Grant 2025-APF-KFZD-30). International Scientific and Technological Cooperation Projects of Sichuan Province, (Grant 2025YFHZ0148). Central Government Guides Local Projects of China, (Grant 2025ZYDF105).

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

1. N. Thuerey, K. Weißenow, L. Prantl, and X. Hu, "Deep Learning Methods for Reynolds-Averaged Navier–Stokes Simulations of Airfoil Flows," *AIAA Journal* 58, no. 1 (2020): 25–36, https://doi.org/10.2514/1.J058291.

2. R. Vinuesa and S. L. Brunton, "Enhancing Computational Fluid Dynamics With Machine Learning," *Nature Computational Science* 2, no. 6 (2022): 358–366, https://doi.org/10.1038/s43588-022-00264-7.

3. Jh Lin, Ss Chen, Jp Li, et al., "A Velocity-Prediction Lightweight Diffusion Model for Three-Dimensional Aircraft Aerodynamic Inverse Design," *Engineering Applications of Artificial Intelligence* 163 (2026): 113125, https://doi.org/10.1016/j.engappai.2025.113125.

4. Li Yg, Ss Chen, Sy Jin, Jh Lin, Ml Jia, and Qf Jiang, "Mixed Conditional Generative Adversarial Networks for High-Dimensional Aircraft Aerodynamic Stealth Shape Optimization," *Engineering Applications of Artificial Intelligence* 173 (2026): 114483, https://doi.org/10.1016/j.engappai.2026.114483.

5. Jh Lin, Ss Chen, Sy Jin, Qf Jiang, Ml Jia, and D. Li, "Multi-Point Aerodynamic Inverse Design of Flying Wing Using the Conditional Diffusion Model," *Physics of Fluids* 37, no. 7 (2025): 077116, https://doi.org/10.1063/5.0271723.

6. T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. Battaglia, *Learning Mesh-Based Simulation With Graph Networks* (OpenReview.net, 2020).

7. D. Roznowicz, G. Stabile, N. Demo, D. Fransos, and G. Rozza, "Large-Scale Graph-Machine-Learning Surrogate Models for 3D-Flowfield Prediction in External Aerodynamics," *Advanced Modeling and Simulation in Engineering Sciences* 11, no. 1 (2024): 6, https://doi.org/10.1186/s40323-024-00259-1.

8. R. F. Miotto and W. R. Wolf, "Flow Imaging as an Alternative to Non-Intrusive Measurements and Surrogate Models Through Vision Transformers and Convolutional Neural Networks," *Physics of Fluids* 35, no. 4 (2023): 045143, https://doi.org/10.1063/5.0144700.

9. X. Wang, S. Zhu, Y. Guo, et al., "TransFlowNet: A Physics-Constrained Transformer Framework For Spatio-Temporal Super-Resolution of Flow Simulations," *Journal of Computational Science* 65 (2022): 101906, https://doi.org/10.1016/j.jocs.2022.101906.

10. X. Guo, W. Li, and F. Iorio, "Convolutional Neural Networks for Steady Flow Approximation," in *Proceedings of the 22nd ACM SIGKDD International Conference*, (2016).

11. V. Sekar, Q. Jiang, C. Shu, and B. C. Khoo, "Fast Flow Field Prediction Over Airfoils Using Deep Learning Approach," *Physics of Fluids* 31, no. 5 (2019): 057103, https://doi.org/10.1063/1.5094943.

12. X. Hui, J. Bai, H. Wang, and Y. Zhang, "Fast Pressure Distribution Prediction of Airfoils Using Deep Learning," *Aerospace Science and Technology* 105 (2020): 105949, https://doi.org/10.1016/j.ast.2020.105949.

13. S. L. Brunton, B. R. Noack, and P. Koumoutsakos, "Machine Learning for Fluid Mechanics," *Annual Review of Fluid Mechanics* 52, no. 1 (2020): 477–508, https://doi.org/10.1146/annurev-fluid-010719-060214.

14. K. Duraisamy, G. Iaccarino, and H. Xiao, "Turbulence Modeling in the Age of Data," *Annual Review of Fluid Mechanics* 51, no. 1 (2019): 357–377, https://doi.org/10.1146/annurev-fluid-010518-040547.

15. Y. Shen, W. Huang, Zg Wang, Df Xu, and C. Y. Liu, "A Deep Learning Framework for Aerodynamic Pressure Prediction on General Three-Dimensional Configurations," *Physics of Fluids* 35, no. 10 (2023): 107111, https://doi.org/10.1063/5.0172437.

16. F. Xiong, L. Zhang, X. Hu, and C. Ren, "A Point Cloud Deep Neural Network Metamodel Method for Aerodynamic Prediction," *Chinese Journal of Aeronautics* 36, no. 4 (2023): 92–103, https://doi.org/10.1016/j.cja.2022.11.025.

17. S. Wang, X. Zhang, Y. Tie, W. Xia, and Y. Jiang, "Aerodynamic Performance Prediction of 3D Aircraft Based on Point Cloud Deep Learning," *Journal of Computational Design and Engineering* 12, no. 5 (2025): 175–190, https://doi.org/10.1093/jcde/qwaf051.

18. M. Elrefaie, A. Dai, and F. Ahmed, "DrivAerNet: A Parametric Car Dataset for Data-Driven Aerodynamic Design and Prediction," *Journal of Mechanical Design* 147, no. 4 (2025): 041712, https://doi.org/10.1115/1.4068104.

19. H. Zhang, Y. Shen, W. Huang, Z. Xie, and Yb Niu, "Deep Transfer Learning For Three-Dimensional Aerodynamic Pressure Prediction Under Data Scarcity," *Theoretical and Applied Mechanics Letters* 15, no. 2 (2025): 100571, https://doi.org/10.1016/j.taml.2025.100571.

20. Z. Zhu, G. Zhao, and Q. Zhao, "Fast and High-Precision Compressible Flowfield Inference Method of Transonic Airfoils Based on Attention Unet," *Physics of Fluids* 36, no. 3 (2024): 036111, https://doi.org/10.1063/5.0188550.

21. K. Zuo, Z. Ye, W. Zhang, X. Yuan, and L. Zhu, "Fast Aerodynamics Prediction of Laminar Airfoils Based on Deep Attention Network," *Physics of Fluids* 35, no. 3 (2023): 037127, https://doi.org/10.1063/5.0140545.

22. M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges," *arXiv preprint arXiv:2104.13478* (2021).

23. P. W. Battaglia, J. B. Hamrick, V. Bapst, et al., "Relational Inductive Biases, Deep Learning, and Graph Networks," *arXiv preprint arXiv:1806.01261* (2018).

24. A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. Battaglia, *Learning to Simulate Complex Physics With Graph Networks* (PMLR, 2020), 8459–8468.

25. F. D. A. Belbute-Peres, T. Economon, and Z. Kolter, *Combining Differentiable PDE Solvers and Graph Neural Networks for Fluid Flow Prediction* (PMLR, 2020), 2402–2411.

26. D. Hines and P. Bekemeyer, "Graph Neural Networks for the Prediction of Aircraft Surface Pressure Distributions," *Aerospace Science and Technology* 137 (2023): 108268, https://doi.org/10.1016/j.ast.2023.108268.

27. C. Ouyang, H. Liu, T. Ma, S. Xu, and X. Xi, "A Physics-Integrated Attention Model With Dynamic Weighting for Rapid and Precise Turbulent Flow Prediction in Variable Geometries," *Physics of Fluids* 37, no. 8 (2025): 085103, https://doi.org/10.1063/5.0278647.

28. A. Hemmasian and F. A. Barati, "Reduced-Order Modeling of Fluid Flows With Transformers," *Physics of Fluids* 35, no. 5 (2023), https://doi.org/10.1063/5.0151515.

29. Q. Cui, M. Zhang, M. Xiao, and G. Ni, "Transformer Based Deep Learning Accelerated Numerical Simulation for Incompressible Flow," *Physics of Fluids* 36, no. 12 (2024): 127152, https://doi.org/10.1063/5.0243174.

30. X. Hu, J. Zhang, K. Yan, T. Wan, and X. Zheng, *Physics-Informed Transformer for Efficient Fluid Dynamics Predictions* (Springer, 2025), 356–368.

31. Y. Fan, C. Wei, J. C. Wong, C. C. Ooi, H. Wang, and P. H. Chiu, "Cascaded PhyFormer-SRNet: A Physics-Informed Transformer-Based Framework With Cascaded Refinement for Spatiotemporal Super-Resolution of Turbulent Flows," *Physics of Fluids* 37, no. 7 (2025): 075207, https://doi.org/10.1063/5.0274979.

32. E. Perez, F. Strub, H. De Vries, V. Dumoulin, and A. Courville, "FiLM: Visual Reasoning With a General Conditioning Layer," in *AAAI'18/IAAI'18/EAAI'18* (AAAI Press, 2018).

33. H. De Vries, F. Strub, J. Mary, H. Larochelle, O. Pietquin, and A. C. Courville, "Modulating Early Visual Processing by Language," in *Advances in Neural Information Processing Systems*, vol. 30 (NeurIPS, 2017).

34. T. Li, L. Guo, Z. Yang, et al., "An Automatic Shape-Aware Method for Predicting Heat Flux of Supersonic Aircraft Based on a Deep Learning Approach," *Physics of Fluids* 34, no. 7 (2022): 077103, https://doi.org/10.1063/5.0098341.