

Article

Physics-Aware Enhanced Implicit Neural Representation with Multimodal Feature Fusion for Zero-Shot Super-Resolution of Complex Flow Fields

Hao Wang, Junlin Fan  and Yin Long * 

School of Computer Science and Technology, Southwest University of Science and Technology, Mianyang 621010, China; 15884242948@163.com (H.W.); fanjunlin@mails.swust.edu.cn (J.F.)

* Correspondence: yinnuolooong@163.com

Abstract

This study addresses the insufficient representation of key regions, such as shock waves and boundary layers, in classical implicit neural representations (INRs) for super-resolution reconstruction of complex aerodynamic flow fields. A physics-aware enhanced implicit neural representation (PA-INR) framework with multimodal feature fusion is proposed. In this study, multimodal features refer to data-driven deep flow features and physics-derived features, including pressure gradients, velocity divergence, and vorticity. The method improves the basic INR in two aspects: local interpolation and feature representation. Specifically, a shock-aware weight correction module is designed to generate corrected weights by combining local features with geometric information, thereby reducing the erroneous fusion of physical states across shock waves. In addition, a gradient-information feature expansion module fuses the data-driven and physics-derived features to enhance the model representation of high-gradient regions. Experimental results show that, on the dataset and evaluation metrics used in this study, PA-INR outperforms the comparison methods in terms of overall reconstruction error and shock-position error. These results indicate that the proposed multimodal feature fusion strategy can improve reconstruction accuracy in complex-flow super-resolution tasks and enhance the recovery of details in shock and high-gradient regions.

Keywords: multimodal feature fusion; implicit neural representation; flow-field super-resolution; zero-shot super-resolution; shock-wave reconstruction; aerodynamic flow fields



Academic Editor: Kung-Ming Chung

Received: 26 July 2026

Revised: 22 August 2026

Accepted: 26 August 2026

Published: 7 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

In modern aerospace engineering, accurate and rapid prediction of complex aerodynamic flow fields around aircraft is fundamental to design optimization, performance evaluation, and flight safety. For a long time, computational fluid dynamics (CFD), which numerically solves the Navier–Stokes equations, has been the most reliable research tool in this field because it can provide high-fidelity flow-field details. However, high-fidelity CFD simulations, especially for flow problems involving turbulence, shock waves, and complex geometries [1,2], usually require enormous computational resources and long solution times. This severely limits their application in scenarios that require rapid design iteration and real-time analysis. Therefore, both academia and industry have been actively exploring more efficient prediction paradigms that can replace or assist conventional CFD. Among them, data-driven methods represented by deep learning [3] have shown great potential.

Flow-field super-resolution is a key technique for improving computational efficiency [4]. It aims to reconstruct high-resolution results [5] from low-resolution flow-field data, thereby substantially reducing the computational cost of CFD while maintaining accuracy. Traditional super-resolution methods, such as bicubic interpolation or spline interpolation, are simple and fast. However, they are essentially numerical interpolation methods based on a smoothness assumption, and they cannot recover high-frequency details lost during downsampling [6]. In key regions such as shocks, vortex structures, and boundary layers, their reconstructions are often overly blurred and physically distorted [7].

To overcome the limitations of traditional methods, researchers have introduced deep learning super-resolution techniques from computer vision into flow-field reconstruction [4]. Among these methods, convolutional neural network (CNN)-based models have achieved notable progress. Early studies directly adopted mature image super-resolution architectures such as SRCNN [8] and EDSR [9]. By performing supervised learning on large numbers of LR–HR flow-field pairs, these networks learn an end-to-end nonlinear mapping from low resolution to high resolution. These methods can recover richer flow-field details than traditional interpolation to some extent. Nevertheless, CNN-based frameworks have several inherent limitations. First, they are usually designed for regular Euclidean data, such as pixel grids, and are less adaptable to flow data on unstructured or curvilinear grids. Second, CNN models typically support only fixed-scale super-resolution [10], lacking the flexibility to reconstruct at arbitrary resolutions. Most importantly, these purely data-driven models lack intrinsic constraints from fluid physics [11]. Their reconstructed fields may therefore violate classical physical laws, appearing visually plausible while being physically incorrect.

In this context, implicit neural representations (INRs) [12] provide a new paradigm for flow-field super-resolution. The core idea of INR is to treat the physical field as a continuous function rather than a collection of discrete data [13,14] and to train a neural network that maps physical condition parameters and spatial coordinates to the physical quantity at the queried point [15]. Because the function is continuous over its domain, the trained model can be queried at coordinate sets of arbitrary density, enabling resolution-independent zero-shot super resolution [16]. This property effectively removes the fixed-scale limitation of CNN models and demonstrates substantial application potential. However, the basic INR framework also has significant intrinsic limitations when applied to aerodynamic problems. These limitations are mainly reflected in two aspects. First, for mathematically discontinuous interfaces such as shock waves, the prediction of a basic INR relies on local interpolation. Geometry-distance-based weighted averaging simply mixes two distinct physical states on the two sides of a shock, leading to severe broadening and blurring of the predicted shock surface. Second, in near-wall boundary layers, flow-field gradients determine key engineering quantities such as friction drag and heat flux. During layer-by-layer abstraction, standard deep networks often smooth or lose these fine but crucial low-level gradient features. Any inaccurate reconstruction of near-wall gradients will directly lead to failure in aircraft performance prediction.

To compensate for the shortcomings of the basic INR in complex flow reconstruction, this study constructs a physics-aware enhanced implicit neural representation framework (PA-INR) [17,18]. The framework focuses on two improvements: reducing cross-discontinuity mixing caused by local geometric interpolation near shocks and strengthening key information in high-gradient regions that is easily weakened during deep feature extraction. Accordingly, this study presents a shock-aware weight correction module and a gradient-information feature expansion module. The former improves local weight allocation, whereas the latter enhances feature representation in high-gradient regions.

Based on these designs, this study aims to improve the reconstruction accuracy of complex aerodynamic flow fields in super-resolution tasks.

2. Related Background

2.1. Flow-Field Super-Resolution

Flow-field super-resolution refers to the recovery and reconstruction of high-resolution flow fields from low-resolution flow-field data [4]. Its goal is to reduce dependence on expensive numerical simulations while preserving the main physical features of the flow as much as possible, thereby improving the efficiency of aerodynamic analysis and design iteration. Early studies mostly adopted traditional reconstruction methods such as bicubic interpolation. However, these methods are essentially numerical approximations based on a smoothness assumption and therefore have difficulty accurately representing high-frequency details contained in regions such as shocks and vortical structures [5–7].

With the development of deep learning, researchers have introduced CNNs into flow-field super-resolution tasks [8] to learn the nonlinear mapping from low-resolution to high-resolution flow fields. These methods improve detail recovery compared with traditional interpolation, but they still have limitations. First, CNNs are generally more suitable for regular grids and have limited adaptability to unstructured grids commonly used in aerodynamics [11,19]. Second, these methods are mostly built for fixed upscaling factors [9,10] and lack the flexibility for arbitrary-resolution reconstruction. Third, purely data-driven frameworks do not explicitly describe the physics of fluid flow and thus may still suffer from insufficient physical consistency in complex-flow regions [17,18].

2.2. Implicit Neural Representations

Implicit neural representations provide a new approach to physical-field modeling that differs from traditional discrete-grid representations. The basic idea is to regard the physical field as a function defined over a continuous space and to approximate this function using a neural network [12,20]. In aerodynamic problems, INR attempts to establish a mapping between physical-condition parameters and spatial locations [21–23], thereby directly predicting the flow variable at any coordinate [16]. Compared with traditional methods that rely on fixed grid resolution, this continuous representation allows the trained model to query arbitrary spatial coordinates and achieve resolution-independent flow-field reconstruction [13,14]. Therefore, it shows considerable potential for zero-shot super-resolution tasks.

Nevertheless, the basic INR still has limitations in aerodynamic scenarios, especially when handling strongly discontinuous or high-gradient regions such as shocks and boundary layers. Its reconstruction results often fail to sufficiently preserve key physical features. Against this background, incorporating physical priors into implicit representations has become an important direction for subsequent methodological improvement. It should be noted that physics-aware representation strategies are different from physics-constrained learning approaches that explicitly impose governing equations or PDE residuals as optimization constraints. In this work, PA-INR introduces physical knowledge through gradient-based feature enhancement and adaptive weight correction, aiming to improve physical consistency while maintaining the flexibility of data-driven implicit representations.

3. Physics-Aware Enhanced Implicit Neural Representation Method

Figure 1 presents the overall architecture of PA-INR and the specific implementations of the shock-aware weight correction module and the gradient-information feature expansion module. It should be noted that the mesh-like structures shown in the schematic figures are only used for illustrating the feature extraction and reconstruction procedures

of PA-INR. They do not represent the actual CFD computational grids used in this study. The original CFD solutions are obtained on hybrid grids and are subsequently converted into regular grids with resolutions ranging from 64×64 to 256×256 for model training and evaluation.




Figure 1. Overall architecture of the PA-INR framework.

3.1. Continuous Flow-Field Reconstruction Based on Implicit Neural Representation

Implicit neural representations model a physical field as a function defined over a continuous spatial domain [12,20], and their core idea is to establish a mapping between physical-condition parameters and spatial coordinates through a neural network [21–23]. In aerodynamics, our objective is to learn a continuous mapping from flow conditions and spatial coordinates to physical quantities. This mapping maps a physical-condition vector $\mu \in \mathbb{R}^{d_\mu}$, which describes the flight state (e.g., Mach number Ma and angle of attack α), to a continuous solution field $u(x; \mu)$ defined over the spatial domain $\Omega \subset \mathbb{R}^d$. INR approximates this continuous function by training a neural network f_θ , thereby constructing a resolution-independent flow representation. The network takes an arbitrary spatial coordinate vector x as input, conditions it on the flow-condition vector μ , and outputs the physical quantity s at that coordinate, such as pressure or velocity components. Because f_θ is continuous over its domain, the trained model can be queried on coordinate sets of arbitrary density [16] to reconstruct flow fields at any target resolution; this forms the theoretical basis of zero-shot super-resolution using INR [13].

To enable a single INR model to represent diverse flow fields determined by different physical conditions μ , we adopt a conditional architecture based on implicit codes. This architecture consists mainly of a hypernetwork h_ψ [15] that encodes flow conditions and a multilayer perceptron (MLP) decoder g_ϕ that decodes spatial coordinates. First, the hypernetwork h_ψ maps the input low-dimensional physical-condition vector μ to a high-dimensional implicit feature code, or latent vector, $z \in \mathbb{R}^{d_z}$, i.e., $z = h_\psi(\mu)$. This latent vector z can be regarded as a compact and abstract mathematical description of the entire flow field under the given condition. The MLP decoder g_ϕ then receives the spatial coordinate vector x and the latent vector z as inputs and outputs the final prediction of the physical quantity. To overcome the spectral bias of standard MLPs when learning high-frequency signals, we first apply a fixed Fourier feature mapping $\gamma(\cdot)$ [20] to the input coordinate x , projecting it into a higher-dimensional feature space to improve the network's

ability to represent flow-field details. Accordingly, the overall INR prediction process can be formulated as follows:

$$s(\mathbf{x}; \boldsymbol{\mu}) \approx g_\phi(\gamma(\mathbf{x}), \mathbf{z}), \quad \text{where } \mathbf{z} = h_\psi(\boldsymbol{\mu}), \quad (1)$$

where the MLP decoder g_ϕ consists of alternating linear layers and nonlinear activation functions σ , such as ReLU, and its information propagation can be written as follows:

$$\mathbf{h}_0 = \text{Concat}(\gamma(\mathbf{x}), \mathbf{z}), \quad \mathbf{h}_{\ell+1} = \sigma(\mathbf{W}_\ell \mathbf{h}_\ell + \mathbf{b}_\ell), \quad g_\phi(\gamma(\mathbf{x}), \mathbf{z}) = \mathbf{W}_L \mathbf{h}_L + \mathbf{b}_L, \quad (2)$$

where $\mathbf{h}_\ell$ denotes the hidden-state vector of the ℓ -th layer, $\mathbf{W}_\ell$ is the weight matrix, and $\mathbf{b}_\ell$ is the bias vector of that layer.

During training, we use a training dataset $\mathcal{D} = \{(\boldsymbol{\mu}_i, \mathbf{U}_{\text{LR}}^{(i)})\}_{i=1}^M$ containing M samples, where each sample includes a physical condition $\boldsymbol{\mu}_i$ and its corresponding CFD ground truth $\mathbf{U}_{\text{LR}}^{(i)}$ solved on the low-resolution (LR) grid $\mathcal{M}_{\text{LR}}^{(i)}$. The LR grid consists of N_{LR} node coordinate vectors $\{\mathbf{x}_j\}_{j=1}^{N_{\text{LR}}}$. We optimize the network parameters $\theta = \{\psi, \phi\}$ by minimizing the mean squared error (MSE) between the model predictions and the ground truth at these LR sampling points. The loss function $\mathcal{L}$ is defined as follows:

$$\mathcal{L}(\theta) = \frac{1}{MN_{\text{LR}}} \sum_{i=1}^M \sum_{j=1}^{N_{\text{LR}}} \|g_\phi(\gamma(\mathbf{x}_j), h_\psi(\boldsymbol{\mu}_i)) - u_{\text{LR}}^{(i)}(\mathbf{x}_j)\|_2^2. \quad (3)$$

Through this training strategy, the model learns a continuous representation of the flow field only from low-cost LR data. However, when this basic INR framework, which is built on continuous function approximation and interpolation, faces the strong discontinuities and high-gradient phenomena characteristic of aerodynamics [1,24,25], its intrinsic limitations become apparent. In particular, when a local interpolation strategy is used for prediction, for example, the prediction at a high-resolution query point $\mathbf{x}_q$ is determined by the four nearest feature vectors $\{\mathbf{z}_t\}_{t \in \{00,01,10,11\}}$ in the LR feature grid, i.e., $s(\mathbf{x}_q) = \sum_t w_t g_\phi(\gamma(\mathbf{x}_q - \mathbf{v}_t), \mathbf{z}_t)$, where the weights w_t are determined by bilinear interpolation area ratios and $\mathbf{v}_t$ is the coordinate of the feature vector $\mathbf{z}_t$. When the query point $\mathbf{x}_q$ lies near a shock wave, this purely geometry-distance-based weighted average incorrectly blends the physically distinct states on the two sides of the shock [25], resulting in a severely broadened and blurred shock surface that cannot represent its true discontinuous physical form. Meanwhile, the loss of details in high-gradient regions such as boundary layers during deep feature extraction further weakens the physical fidelity of the model predictions. These limitations motivate us to develop enhanced modules that incorporate physical prior knowledge more explicitly into the INR framework [17,18].

3.2. Shock-Aware Weight Correction Module

To effectively handle local physical phenomena in the flow field, especially accurate reconstruction near discontinuities such as shock waves, we extend the INR model from a framework that relies on a single global implicit code to one that can perceive and respond to local flow-field features. To this end, we first introduce a backbone encoder [26], such as EDSR [9] or RDN [10], which transforms the input LR flow field $\mathbf{U}_{\text{LR}}$ into a spatialized LR feature grid $\mathbf{Z}_{\text{LR}} \in \mathbb{R}^{H_{\text{LR}} \times W_{\text{LR}} \times C}$. Each element $\mathbf{z}_{ij}$ in this grid is a deep feature vector containing flow information from the neighborhood of its corresponding spatial location. Based on this feature grid, the prediction at a high-resolution (HR) query point is determined by its surrounding local features rather than by a global code.

Within this localized framework, a basic prediction strategy is local ensemble, which has been proven effective in image super-resolution models such as LIIF [16]. Specifically,

for an arbitrary query coordinate $\mathbf{x}_q$ on the HR grid, we first map it to the continuous coordinate system of the LR feature grid. Then, we identify the four nearest feature vectors surrounding $\mathbf{x}_q$, denoted by $\{\mathbf{z}_t\}_{t \in \{00,01,10,11\}}$, and their corresponding coordinates in the feature grid, denoted by $\{\mathbf{v}_t\}_{t \in \{00,01,10,11\}}$. The local ensemble strategy computes the final prediction through bilinear interpolation [16], with weights determined entirely by the relative geometric positions between the query point and the four neighboring feature points. Suppose that $\mathbf{x}_q$ has local coordinates $(\Delta x, \Delta y)$ inside the unit cell enclosed by $\{\mathbf{v}_t\}$, where $\Delta x, \Delta y \in [0, 1]$. Then, the geometric interpolation weights w_t of the four neighbors can be expressed as:

$$\begin{aligned} w_{00} &= (1 - \Delta x)(1 - \Delta y), \\ w_{01} &= (1 - \Delta x)\Delta y, \\ w_{10} &= \Delta x(1 - \Delta y), \\ w_{11} &= \Delta x \Delta y. \end{aligned} \quad (4)$$

At this point, the purely geometry-based interpolation result can be obtained by the following equation:

$$s_{\text{geom}}(\mathbf{x}_q) = \sum_{t \in \{00,01,10,11\}} w_t g_\phi(\gamma(\mathbf{x}_q - \mathbf{v}_t), \mathbf{z}_t). \quad (5)$$

This equation reveals the source of the problem. In aerodynamics, a shock wave is a mathematical discontinuity in the flow state [1,24]. Although the physical states on the two sides of a shock are spatially adjacent, they are almost decoupled physically. The purely geometric weighting method inevitably averages or smooths this discontinuous physical jump [25], causing the predicted shock surface to become blurred, which is physically incorrect.

To overcome this defect, we design and introduce a shock-aware weight correction module $\mathcal{P}_\xi$, whose core function is to learn a weight allocation mechanism that perceives the local physical structure of the flow rather than the geometric structure alone. The module is a lightweight neural network composed of several fully connected layers, with parameters denoted by ξ . Unlike simple approaches that take only geometric weights as inputs, we concatenate the initial geometric weights $\{w_t\}$ with their corresponding four local feature vectors $\{\mathbf{z}_t\}$ to form a richer input vector $\mathbf{i}_q$, enabling the module to make physically informed judgments:

$$\mathbf{i}_q = \text{Concat}(w_{00}, w_{01}, w_{10}, w_{11}, \mathbf{z}_{00}, \mathbf{z}_{01}, \mathbf{z}_{10}, \mathbf{z}_{11}). \quad (6)$$

By including the local feature vectors themselves as part of the input, the weight correction module $\mathcal{P}_\xi$ can directly observe the flow-feature distribution around the query point. During end-to-end training, $\mathcal{P}_\xi$ learns to identify feature patterns that imply the presence of a shock, such as sharp and directional differences between adjacent feature vectors. When such patterns are detected, it outputs a set of corrected weights $\{w'_t\}$ related to the physical structure. The last layer of the module uses a Softmax activation function to ensure that the output weights sum to 1, maintaining their physical meaning as a valid probability distribution:

$$[w'_{00}, w'_{01}, w'_{10}, w'_{11}]^\top = \text{Softmax}(\mathcal{P}_\xi(\mathbf{i}_q)). \quad (7)$$

To retain the stability advantage of geometric interpolation in smooth regions while introducing necessary physical corrections in discontinuous regions such as shock waves, we use a weighted fusion strategy to obtain the final weights $\{w''_t\}$. This strategy bal-

ances the original geometric weights and the learned physics-corrected weights through a hyperparameter weight factor $\lambda_c \in [0, 1]$:

$$w_t'' = (1 - \lambda_c)w_t + \lambda_c w_t', \quad t \in \{00, 01, 10, 11\}. \quad (8)$$

Finally, after the physics-aware weight correction, the physical quantity predicted at the HR query point x_q , denoted as $s_{\text{final}}(x_q)$, is computed as follows:

$$s_{\text{final}}(x_q) = \sum_{t \in \{00, 01, 10, 11\}} w_t'' g_\phi(\gamma(x_q - v_t), z_t). \quad (9)$$

With this design (Figure 2), the model can retain the stability of geometric interpolation in smooth regions while reducing errors caused by cross-discontinuity mixing near shock waves, thereby improving reconstruction in local key regions. The effectiveness of the shock-aware correction mechanism originates from the difference between geometric proximity and physical similarity. Although neighboring feature points are close in the spatial domain, they may represent substantially different flow states when separated by a shock discontinuity. Therefore, directly applying distance-based interpolation may introduce physically inconsistent feature mixing. By incorporating local feature information into the weight prediction process, the proposed module enables adaptive interpolation according to local flow characteristics rather than spatial distance alone.




Figure 2. Schematic diagram of the shock-aware weight correction module.

3.3. Gradient-Information Feature Expansion Module

In viscous fluid dynamics, the near-wall boundary layer is a core region that determines key engineering quantities such as aircraft skin-friction drag and surface heat flux [2]. These quantities are physically related to the flow-field gradients at the wall. For example, the wall shear stress τ_w , which further determines the friction coefficient C_f , is proportional to the wall-normal velocity gradient, whereas the wall heat flux q_w is proportional to the wall-normal temperature gradient. These relationships can be formulated as follows:

$$\tau_w = \mu_f \left. \frac{\partial u_t}{\partial y} \right|_{y=0}, \quad q_w = -k \left. \frac{\partial T}{\partial y} \right|_{y=0}, \quad (10)$$

where μ_f is the dynamic viscosity, k is the thermal conductivity, u_t is the tangential velocity, T is the temperature, and y is the wall-normal coordinate. This relationship indicates that any

inaccurate reconstruction of near-wall gradients will lead to errors in key aircraft-performance quantities. However, standard deep neural networks, such as CNNs, often smooth or lose these fine but crucial low-level gradient details in the original LR input field during the layer-by-layer abstraction of high-level semantic features. This information loss affects the reconstruction accuracy of INR models in high-gradient regions, especially in engineering applications that require accurate prediction of aerodynamic forces and heat transfer.

To address this issue, we design and introduce a gradient-information feature expansion module. Its core task is to explicitly calculate and preserve key gradient information from the input flow field before deep feature extraction [17,18] and to feed it into subsequent networks as explicitly constructed physical features. Specifically, for a given multichannel LR input flow field $\mathbf{U}_{\text{LR}}$, which contains physical quantities such as pressure p and velocity components (u, v) , we first apply a set of physics-selected numerical differential operators $\mathcal{G}$ [11] to generate a multichannel gradient feature map $\mathbf{Z}_{\text{grad}} = \mathcal{G}(\mathbf{U}_{\text{LR}})$. These operators are not chosen arbitrarily; rather, they are selected to capture physical quantities that strongly affect flow structures. We use numerical schemes such as second-order central differences to approximate spatial derivatives and generate the following key gradient feature channels:

1. Pressure-gradient magnitude G_p : a robust indicator of shocks, expansion waves, and strong adverse-pressure-gradient regions:

$$G_p(x, y) = \sqrt{\left(\frac{\partial p}{\partial x}\right)^2 + \left(\frac{\partial p}{\partial y}\right)^2}. \quad (11)$$

2. Velocity divergence G_{div} : in compressible flow, it directly measures the volumetric deformation rate of the fluid and serves as a direct indicator for distinguishing compression regions ($G_{\text{div}} < 0$) from expansion regions ($G_{\text{div}} > 0$):

$$G_{\text{div}}(x, y) = \frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}. \quad (12)$$

3. Vorticity magnitude G_{ω} : as the magnitude of the curl of the velocity field, it is twice the angular velocity of a fluid element and is a direct physical quantity for identifying rotational flow structures such as shear layers, boundary layers, and wake vortices:

$$G_{\omega}(x, y) = \left| \frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right|. \quad (13)$$

The calculated gradient feature map $\mathbf{Z}_{\text{grad}} = [G_p, G_{\text{div}}, G_{\omega}, \dots]$ has the same spatial dimensions $H_{\text{LR}} \times W_{\text{LR}}$ as the original LR input $\mathbf{U}_{\text{LR}}$.

Next, we fuse the explicitly calculated physical gradient information [11] with the deep abstract features extracted by the backbone encoder $\mathcal{E}_{\eta}$ [9,10]. The backbone encoder, parameterized by η , maps the original LR input $\mathbf{U}_{\text{LR}}$ to a high-dimensional deep feature map $\mathbf{Z}_{\text{deep}} = \mathcal{E}_{\eta}(\mathbf{U}_{\text{LR}})$ with dimensions $H_{\text{LR}} \times W_{\text{LR}} \times C_{\text{deep}}$. Our fusion strategy concatenates them along the channel dimension, producing an enhanced feature grid $\mathbf{Z}_{\text{enhanced}}$ that is richer in information and contains both high-level semantic features and low-level physical gradients:

$$\mathbf{Z}_{\text{enhanced}} = \text{Concat}(\mathbf{Z}_{\text{deep}}, \mathbf{Z}_{\text{grad}}). \quad (14)$$

The enhanced feature grid has dimensions $H_{\text{LR}} \times W_{\text{LR}} \times (C_{\text{deep}} + C_{\text{grad}})$. It replaces the original deep feature map and serves as the data basis for the local ensemble and weight correction module described in the preceding section. Therefore, the final local

feature vectors $\{z_t\}$ used for prediction are sampled from $\mathbf{Z}_{\text{enhanced}}$. The construction of the enhanced feature grid is illustrated in Figure 3.

This design improves the performance of the entire INR framework. When the final MLP decoder g_ϕ receives the local feature vector $z_{\text{enhanced},t}$ from the enhanced feature grid, it no longer relies only on deep features to infer high-gradient regions indirectly [17,18]. Instead, explicit information about the magnitude and location of gradients has already been encoded directly in the input features. This effectively reduces the learning difficulty of the MLP and allows more network capacity to be used for accurately fitting the complex functional forms in these gradient regions [12,20] rather than for detecting whether such regions exist. From a physical perspective, the gradient-information feature expansion provides a complementary representation to the learned deep features. The deep feature encoder mainly captures abstract spatial patterns from the flow field, whereas the explicitly calculated gradients provide direct descriptions of local variation intensity. This combination is particularly important for aerodynamic fields containing shocks, boundary layers, and other high-gradient structures, where purely data-driven feature extraction may suppress small-scale but physically important variations.

Finally, the prediction formula weighted by the corrected weights, $s_{\text{final}}(x_q) = \sum_t w_t'' g_\phi(\gamma(x_q - v_t), z_{\text{enhanced},t})$, can generate high-resolution flow fields with sharper and more accurate profiles in key regions such as boundary layers and shocks, thereby providing more reliable reconstructed results for subsequent engineering analysis.




Figure 3. Schematic diagram of the gradient-information feature expansion module. (a) Gradient Information Feature Expansion Module; (b) Physics-Aware Weight Correction Module.

4. Experimental Design and Result Analysis

To quantitatively evaluate the effectiveness of the proposed physics-aware enhanced implicit neural representation (PA-INR) framework, a series of validation experiments are designed around transonic airfoil flow. This section describes the dataset, model settings, training details, and evaluation metrics used in the experiments.

4.1. Dataset and Data Resolution

(1) Dataset

To evaluate the proposed model comprehensively and rigorously, we use a representative and publicly available benchmark dataset for aerodynamic surrogate modeling [3,27,28].

2D Transonic Airfoil Dataset: This dataset contains a large number of Reynolds-averaged Navier–Stokes (RANS) numerical solutions around the classic transonic airfoil RAE2822 [28]. It consists of 1200 simulation samples with a fixed geometry but a wide range of flow conditions: the Mach number ranges from 0.3 to 0.9, and the angle of attack ranges from 0° to 9° . The Reynolds number is fixed at 6.5 million for all simulations, and the Spalart–Allmaras turbulence model [2] is used. Each simulation is performed on a hybrid grid with 27,499 cells. The value of this dataset lies in its coverage of the full flow-regime transition from weakly compressible to strongly transonic flow, including many samples with clear shock structures [3]. This makes it an ideal platform for testing the model’s ability to capture high-gradient and strong-discontinuity phenomena and for conducting detailed evaluation of approximate discretization invariance and continuous super-resolution performance [27].

(2) Data Resolution

The low-resolution (LR) training inputs are generated from the HR data by average pooling with a stride of 4, resulting in a resolution of 64×64 . The training label for each sample is the pressure field at 128×128 resolution. To evaluate model performance at arbitrary resolutions, the target resolutions are set to 128×128 , 192×192 , and 256×256 according to the experimental requirements.

4.2. Experimental Settings

All models are implemented using the PyTorch 2.7.0 deep learning framework [29]. Training is performed with the Adam optimizer [30], with an initial learning rate of 1×10^{-4} and cosine annealing [31] for learning rate decay. The dataset is divided into training, validation, and test sets in proportions of 70%, 10%, and 20%, respectively. The model loss function is the mean squared error (MSE) between the predicted values and the CFD ground truth. In the super-resolution experiments, models are trained using LR inputs and tested on grids with different target resolutions to assess their arbitrary-resolution reconstruction capability.

(1) Comparison Methods

To verify the effectiveness of the proposed method, PA-INR is compared with the following three baseline models:

Aero-Nef [27]: This is an important comparison method in this study and represents a class of implicit neural representation methods using a global conditioning strategy. Comparison with Aero-Nef makes it possible to analyze the improvement of the proposed method in reconstructing local high-frequency details.

Graph neural network [32]: A multiscale graph neural network model [19,33] (GUNET) is selected for comparison. This class of methods is an important technical route for processing unstructured grids and is used to compare the performance of the INR framework and graph neural networks in super-resolution tasks.

Standard INR: This is a simplified INR model that directly feeds coordinates and global parameters into a standard multilayer perceptron (MLP) without any specially designed processor or decoder. This baseline serves as an ablation study reference for evaluating the improvements brought by the proposed modules over the basic INR framework.

(2) PA-INR Model Settings

Backbone encoder: We use EDSR [9] as the backbone encoder to extract a spatialized deep feature grid $\mathbf{Z}_{\text{deep}}$ from the LR flow input $\mathbf{U}_{\text{LR}}$. The encoder consists of 16 residual blocks [26], and each convolutional layer has 64 feature channels.

Gradient-information feature expansion module: Following Section 3.3, this module first computes three gradient feature maps from the input flow field $\mathbf{U}_{\text{LR}}$: pressure-gradient magnitude G_p , velocity divergence G_{div} , and vorticity magnitude G_ω . These gradient maps

are then concatenated with the deep feature map $\mathbf{Z}_{\text{deep}}$ output by the backbone encoder along the channel dimension, forming an enhanced feature grid $\mathbf{Z}_{\text{enhanced}}$ with dimensions $H_{\text{LR}} \times W_{\text{LR}} \times (C_{\text{deep}} + C_{\text{grad}})$.

Shock-aware weight correction module $\mathcal{P}_{\xi}$: This module is a lightweight MLP with two hidden layers that learns physics-corrected weights. The input layer receives the concatenated geometric weights and local feature vectors, the hidden dimension is 32, and the activation function is ReLU. The output layer uses Softmax activation to ensure weight normalization. The weight-fusion hyperparameter λ_c is determined through a sensitivity analysis on the validation set. The detailed analysis is presented in Section 5. Based on the results, $\lambda_c = 0.8$ is adopted in all subsequent experiments.

MLP decoder: The decoder g_{ϕ} contains four hidden layers, each with 256 units. It receives the relative coordinate encoded by the Fourier feature mapping $\gamma(x)$ and the enhanced local feature vector sampled from $\mathbf{Z}_{\text{enhanced}}$ as inputs. ReLU is also used as the activation function.

(3) Quantitative Metrics:

1. Mean Squared Error (MSE): the primary metric for evaluating overall flow-field reconstruction accuracy.
2. Peak Signal-to-Noise Ratio (PSNR): a standard metric borrowed from image processing to evaluate the quality of the predicted field and its sensitivity to the error distribution.
3. Shock-Position Error [1,24]: this metric evaluates the deviation of the predicted shock position and is an important indicator for measuring the model reconstruction capability for local strong-discontinuity structures such as shocks [25]. It is calculated by comparing the shock positions in the prediction and the CFD ground truth. A smaller error indicates more accurate shock-region reconstruction.

(4) Qualitative Metrics:

1. Flow-field contour comparison: the predicted pressure or Mach-number contours are visualized and compared side by side with the CFD ground truth to intuitively evaluate the model capability to capture complex flow structures such as shocks and vortices.
2. Surface pressure distribution comparison: pressure distribution curves on the upper and lower surfaces of the airfoil are extracted and compared with the CFD ground truth to evaluate the recovery of leading-edge pressure variation, local gradient variation, and the overall pressure-distribution trend.

4.3. Experimental Results and Analysis

(1) Quantitative Evaluation

This section quantitatively evaluates the performance of PA-INR in arbitrary-scale super-resolution tasks. All models are trained on low-resolution inputs from the 2D transonic airfoil dataset, with the corresponding pressure fields at a resolution of 128×128 used as training targets. After training, the models are evaluated on held-out test samples at target resolutions of 128×128 , 192×192 , and 256×256 . The 128×128 test case uses unseen test samples at a resolution employed during training, whereas the 192×192 and 256×256 target resolutions are not used during training. The evaluation metrics include MSE, PSNR, and shock-position error. The first two metrics measure overall reconstruction accuracy, whereas the last evaluates the ability of the model to preserve the location of discontinuity structures. Specifically, the shock position is identified from the pressure field by detecting the region with the strongest pressure variation along the streamwise direction. The detected shock locations from the reconstructed field and the CFD reference field are then compared using the absolute spatial distance. Since all reconstructed fields

are evaluated on regular grids, the shock-position error is reported in pixels corresponding to the target resolution.

Table 1 presents the super-resolution results of all models at the target resolution of 128×128 . PA-INR outperforms the comparison methods on all three metrics, namely MSE, PSNR, and shock-position error, indicating good performance in both overall flow-field reconstruction and shock-region prediction. Compared with Aero-Nef, PA-INR shows a more evident advantage in shock-position error, suggesting that the proposed method is better adapted to reconstructing local high-frequency structures. The metrics of GNN and standard INR are relatively weaker, indicating that their reconstruction accuracy remains limited in the present task.

Table 1. Super-resolution performance comparison from low resolution to a 128×128 grid.

Model	Mean Squared Error (MSE)	Peak Signal-to-Noise Ratio (PSNR)	Shock-Position Error (Pixels)
PA-INR	0.0018	27.16	1.48
Aero-Nef	0.0028	25.53	2.87
GNN	0.0085	20.71	6.54
Standard INR	0.0041	23.87	3.98

When the target resolution is increased to 192×192 , the reconstruction difficulty of all models increases further. The results show that the MSE and shock-position error of PA-INR increase to some extent, but the overall degradation is relatively small, and the model still maintains good performance. In contrast, the errors of the other baselines increase more noticeably, especially the shock-position error of Aero-Nef. This indicates that, under higher-scale super-resolution tasks, the proposed method can still maintain strong local-detail reconstruction capability.

In the high-scale super-resolution test at 256×256 , the performance differences among models become larger. PA-INR still maintains relatively good performance and outperforms the comparison methods on MSE, PSNR, and shock-position error. This indicates that, as the target resolution increases, the performance degradation of the proposed method is relatively small and the model is stable. In comparison, the errors of Aero-Nef, GNN, and standard INR increase more significantly, suggesting that these methods face greater difficulty in recovering details at higher reconstruction resolutions. Overall, the results in Tables 1–3 show that the proposed gradient-information feature expansion and local weight correction mechanisms help improve reconstruction accuracy in complex-flow super-resolution tasks, especially in shock regions.

Table 2. Super-resolution performance comparison from low resolution to a 192×192 grid.

Model	Mean Squared Error (MSE)	Peak Signal-to-Noise Ratio (PSNR)	Shock-Position Error (pixels)
PA-INR	0.0026	26.55	2.12
Aero-Nef	0.0039	24.09	4.15
GNN	0.0121	19.17	8.92
Standard INR	0.0058	22.37	5.66

Table 3. Super-resolution performance comparison from low resolution to a 256×256 grid.

Model	Mean Squared Error (MSE)	Peak Signal-to-Noise Ratio (PSNR)	Shock-Position Error (pixels)
PA-INR	0.0036	25.68	2.91
Aero-Nef	0.0055	22.60	5.98
GNN	0.0173	17.62	12.43
Standard INR	0.0082	20.86	7.81

(2) Ablation Study of PA-INR Components

To investigate the contribution of each component in PA-INR, we conduct a component-level ablation study. The baseline Standard INR is gradually extended by introducing the EDSR backbone, gradient-information feature expansion module, and shock-aware weight correction module. The quantitative results of the ablation study are presented in Table 4.

Table 4. Ablation study of different PA-INR components at a target resolution of 256×256 .

Model	MSE	PSNR	HG-MSE	Shock Error
Standard INR	0.0082	20.86	0.0238	7.81
EDSR-INR	0.0055	22.75	0.0168	5.63
EDSR-INR + Gradient	0.0046	23.85	0.0120	5.05
EDSR-INR + Shock	0.0048	23.58	0.0135	3.86
PA-INR	0.0036	25.68	0.0091	2.91

The ablation results demonstrate that each component contributes to the improvement of PA-INR. Compared with Standard INR, EDSR-INR achieves lower reconstruction error, indicating the effectiveness of localized feature representation. After introducing the gradient-information feature expansion module, the HG-MSE decreases from 0.0168 to 0.0120, demonstrating improved reconstruction capability in high-gradient regions. Meanwhile, adding the shock-aware weight correction module reduces the shock-position error from 5.63 pixels to 3.86 pixels, confirming its effectiveness in mitigating cross-discontinuity interpolation errors. Finally, the complete PA-INR achieves the best performance, indicating that the two proposed modules provide complementary improvements.

(3) Sensitivity Analysis of Weight-Fusion Factor λ_c

To investigate the influence of the weight-fusion factor λ_c on reconstruction performance, we conduct a sensitivity analysis with different values of λ_c . The parameter controls the balance between geometry-based interpolation weights and learned physics-aware correction weights. When $\lambda_c = 0$, the model relies completely on geometric interpolation, while $\lambda_c = 1$ represents fully learned weight correction. The experiments are conducted under identical network structures and training configurations. The quantitative results are presented in Table 5.

The results show that increasing λ_c from 0 to 0.8 gradually improves reconstruction performance. This indicates that the learned physics-aware correction effectively reduces cross-discontinuity interpolation errors near shock regions. However, when λ_c is further increased to 0.9 or 1.0, the performance slightly decreases. This is because completely replacing the geometric interpolation prior with learned weights may reduce reconstruction stability in smoother regions. Therefore, $\lambda_c = 0.8$ provides a favorable balance between geometric stability and physics-aware correction capability.

(4) Flow-Field Contours and Error-Distribution Comparison

As shown in Figure 4, under the 7.2° condition, the errors of all models are mainly concentrated near the airfoil leading edge. The error distribution of GNN is broader and

includes a noticeable secondary error region near the trailing edge. Aero-Nef and the standard INR mainly show localized error concentrations near the leading edge, but still retain some residual errors near the trailing edge. In contrast, the high-error region of the proposed method is confined to a very small leading-edge area, and the overall error distribution is more concentrated. This indicates that the proposed method reconstructs the key flow structures more closely to the CFD result under this condition.

Table 5. Sensitivity analysis of the weight-fusion factor λ_c at different target resolutions.

λ_c	MSE	PSNR	Shock Error	Resolution
0.0	0.00531	23.62	5.35	256×256
0.2	0.00492	24.03	4.82	256×256
0.4	0.00448	24.47	4.24	256×256
0.6	0.00398	25.02	3.51	256×256
0.7	0.00373	25.39	3.12	256×256
0.8	0.00360	25.68	2.91	256×256
0.9	0.00370	25.48	3.04	256×256
1.0	0.00396	25.10	3.36	256×256




Figure 4. Static pressure-field reconstruction contours and absolute error distributions at an angle of attack of 7.200° .

As shown in Figure 5, under the 4.17° condition, the errors of all models are mainly concentrated near the airfoil leading edge and the front part of the upper surface, and the overall error level is significantly lower than that under the 7.2° condition. The standard INR shows noticeable error regions on both the upper and lower sides of the leading edge, whereas GNN exhibits stronger error diffusion. Aero-Nef has a relatively concentrated error distribution, but it still retains a clear local error band above the leading edge. In contrast, the high-error region of the proposed method is mainly confined to a small neighborhood of the leading edge, with weak errors over the remaining airfoil surface. This indicates that the proposed method better preserves the reconstruction accuracy of local flow structures under this condition.




Figure 5. Static pressure-field reconstruction contours and absolute error distributions at an angle of attack of 4.170° .

As shown in Figure 6, under the 8.00° condition, the errors of all models are mainly concentrated near the airfoil leading edge, indicating that this region is the dominant source of error in the static-pressure-field reconstruction for the current condition. The error diffusion range of GNN is relatively large, while Aero-Nef and the standard INR mainly show leading-edge local error concentrations. The high-error region of the proposed method is further restricted to a smaller area, and the overall error distribution is more concentrated. This suggests that the proposed method can still control local reconstruction errors well under larger angles of attack.

As shown in Figure 7, under the 3.45° condition, the errors of all models are mainly concentrated near the airfoil leading edge and the front portions of the upper and lower surfaces. GNN and the standard INR have relatively larger error diffusion ranges and show more evident local propagation characteristics. The error distributions of Aero-Nef and the proposed method are relatively more concentrated and mainly limited to the leading-edge neighborhood. Overall, the proposed method can control the errors within local key regions and exhibits stable reconstruction results.




Figure 6. Static pressure-field reconstruction contours and absolute error distributions at an angle of attack of 8.000° .

Combining the reconstructed contours and error distributions under different angles of attack shows that the main errors of all models are mostly concentrated near the airfoil leading edge and its neighboring region. This indicates that the leading-edge region is consistently the most sensitive and difficult region to recover accurately in static-pressure-field super-resolution reconstruction. Although the error peak values and diffusion ranges vary with flow condition, PA-INR generally exhibits stronger error-localization capability under all conditions. Its high-error regions are mainly restricted to small neighborhoods of the leading edge, while the errors in the middle and rear parts of the airfoil and in the far field are generally weak. In contrast, GNN and the standard INR are more prone to error diffusion. Aero-Nef can maintain relatively concentrated error distributions under some conditions, but still shows clear local error bands near the leading edge. The reconstructed

contours further indicate that PA-INR is closer to the CFD ground truth in terms of shock position, local gradient variations, and pressure recovery near the airfoil surface. Overall, the proposed method demonstrates good reconstruction stability under different angles of attack. It maintains overall reconstruction accuracy while more effectively controlling local errors in key regions. This suggests that the gradient-information feature expansion module and the shock-aware weight correction module play positive roles in detail recovery and error suppression for complex-flow super-resolution tasks.




Figure 7. Static pressure-field reconstruction contours and absolute error distributions at an angle of attack of 3.450° .

(5) Surface Pressure Distribution Comparison

As shown in Figure 8, under the $\alpha = 7.200^\circ$ condition, the surface pressure distribution of the airfoil changes noticeably, and a strong pressure jump exists near the leading edge. The curve comparison shows that PA-INR agrees more closely with the CFD ground truth, especially near the leading-edge peak and in the local variation region downstream, where the prediction is closer to the true distribution. In contrast, GNN shows more evident local deviations, while INR and Aero-Nef can capture the overall trend but still have errors near local turning points. This indicates that, under conditions with strong pressure variation, PA-INR has stronger recovery capability in local key regions.

As shown in Figure 9, under the $\alpha = 4.170^\circ$ condition, all models show high overall agreement with the CFD ground truth, indicating that the overall reconstruction difficulty of the surface pressure distribution is relatively low for this condition. However, local comparisons still reveal differences among models near the leading edge and in transition regions. Overall, PA-INR is closer to the CFD result in terms of the leading-edge peak position, curve-slope variation, and smooth transition of the pressure distribution on the upper and lower surfaces. It therefore shows good fitting performance and stability. In comparison, the other models can reflect the overall pattern of the airfoil surface pressure distribution, but slight deviations remain in local regions with rapid changes. This shows that PA-INR retains good detail-preservation capability even under relatively smooth overall variations.




Figure 8. Surface pressure distribution comparison at an angle of attack of 7.200° .




Figure 9. Surface pressure distribution comparison at an angle of attack of 4.170° .

As shown in Figure 10, under the $\alpha = 8.000^\circ$ condition, the pressure variation near the airfoil leading edge is further intensified, imposing higher requirements on the model's

ability to recover local strong-gradient information. The results show that PA-INR still tracks the overall shape of the CFD curve well and achieves high fitting accuracy in both the leading-edge strong-variation region and the subsequent pressure-recovery process. In contrast, the other models can reflect the overall trend, but deviations can still be observed near the leading-edge extrema and in local transition regions.

As shown in Figure 11, under the $\alpha = 3.450^\circ$ condition, the overall surface pressure distribution of the airfoil varies relatively smoothly, but a clear pressure peak and large local gradient still exist near the leading edge, where prediction differences among models are most concentrated. The curve comparison shows that all methods can follow the overall variation trends of the upper and lower airfoil surfaces in the middle and rear chord regions, indicating that the overall reconstruction difficulty of the surface pressure distribution is relatively low under this small-angle-of-attack condition. However, near the leading edge, GNN, INR, and Aero-Nef still show varying degrees of local deviation, especially around pressure extrema and rapid curve-turning locations, where their predictions slightly deviate from the CFD ground truth. In contrast, the proposed method agrees more closely with the CFD curve, accurately recovering the leading-edge peak, the pressure difference between the upper and lower surfaces, and the subsequent pressure recovery along the chord. As the x -coordinate increases, the curves of all models gradually converge, but the proposed method maintains more stable fitting over the full chord length. This indicates that the gradient-information feature expansion module and the physics-aware weight correction mechanism help improve recovery of leading-edge high-gradient regions and surface-pressure details. Overall, under the $\alpha = 3.450^\circ$ condition, the proposed method not only preserves the overall trend of the surface pressure distribution but also shows better reconstruction accuracy and stability in local key regions near the leading edge.




Figure 10. Surface pressure distribution comparison at an angle of attack of 8.000° .

The surface pressure distribution curves under different conditions show that PA-INR performs better overall in recovering the pressure distribution on the upper and lower airfoil surfaces. In particular, in the leading-edge peak region, local rapid-variation regions, and the middle-to-rear pressure recovery region, the proposed method more accurately preserves the curve shape and overall variation trend, making the predictions closer to the CFD ground truth. Combined with the preceding quantitative metrics and flow-field contour analyses, PA-INR can be considered to have advantages in both overall reconstruction accuracy and local key-feature recovery.




Figure 11. Surface pressure distribution comparison at an angle of attack of 3.450° .

(6) Computational Complexity and Efficiency Analysis

To evaluate whether the accuracy improvement of PA-INR introduces excessive computational overhead, we further analyze the computational efficiency of different models. The comparison results are summarized in Table 6, including model parameters, inference latency, and peak GPU memory consumption. The experiments are conducted under identical hardware and software settings to ensure a fair comparison.

Table 6. Accuracy–efficiency comparison of different models at 256×256 resolution. The symbol $\downarrow$ indicates that lower values are better.

Model	MSE $\downarrow$	Shock Error $\downarrow$	Parameters (M)	Inference Time (ms) $\downarrow$	Memory (GB) $\downarrow$
Standard INR	0.0082	7.81	0.56	19.2	1.05
Aero-Nef	0.0055	5.98	1.21	27.2	1.34
GNN	0.0173	12.43	3.84	43.8	2.42
EDSR-INR	0.0055	5.63	1.680	30.7	1.62
EDSR-INR+Grad	0.0046	5.05	1.681	32.1	1.70
EDSR-INR+Shock	0.0048	3.86	1.692	34.2	1.76
PA-INR	0.0036	2.91	1.693	35.7	1.84

Compared with the EDSR-INR baseline, PA-INR achieves substantial improvements in reconstruction accuracy with only moderate additional computational cost. The parameter number increases from 1.680 M to 1.693 M, corresponding to an increase of approximately 0.8%. Meanwhile, the inference time increases from 30.7 ms to 35.7 ms, while the peak inference memory increases from 1.62 GB to 1.84 GB. In contrast, the MSE decreases from 0.0055 to 0.0036, and the shock-position error decreases from 5.63 pixels to 2.91 pixels. These results demonstrate that the proposed physics-aware modules provide significant accuracy improvements while maintaining a favorable computational efficiency trade-off. It is worth noting that the reconstruction capability of PA-INR depends on the quality of the low-resolution input field. The input data should preserve the major physical characteristics of the original flow, such as the global pressure distribution, dominant gradients, and

main discontinuity locations. When important flow structures are completely lost during downsampling, they cannot be recovered solely through implicit representation. Therefore, appropriate selection of the input resolution is essential to maintain sufficient physical information while achieving efficient super-resolution reconstruction.

5. Conclusions

This study proposed a physics-aware enhanced implicit neural representation framework, named PA-INR, for zero-shot super-resolution reconstruction of complex aerodynamic flow fields. The method was designed to address the limitations of classical implicit neural representations in capturing discontinuous and high-gradient flow structures, especially shock waves and boundary-layer-related regions. To this end, a shock-aware weight correction module was introduced to reduce erroneous cross-discontinuity interpolation, and a gradient-information feature expansion module was developed to explicitly incorporate pressure-gradient magnitude, velocity divergence, and vorticity information into the feature representation.

Experiments on the 2D transonic airfoil dataset demonstrate that PA-INR achieves better reconstruction performance than Aero-Nef, GNN, and standard INR under different target resolutions. The proposed method obtains lower mean squared error, higher PSNR, and smaller shock-position error at 128×128 , 192×192 , and 256×256 resolutions. The flow-field contour comparisons and surface pressure distribution analyses further show that PA-INR can more accurately recover local pressure variations, reduce error diffusion near the airfoil leading edge, and better preserve key flow structures compared with the baseline methods.

Overall, the results indicate that introducing physics-aware local weight correction and explicit gradient-feature enhancement can effectively improve the accuracy and stability of implicit neural representations for complex-flow super-resolution tasks. The proposed PA-INR framework provides a useful approach for efficient aerodynamic flow-field reconstruction and has potential value in rapid aerodynamic analysis and design optimization. Future work will investigate the extension of PA-INR to more complex three-dimensional configurations, multiple aerodynamic variables, and broader flow conditions. In addition, incorporating stronger physical constraints or governing equation-related information will be explored to further improve generalization and physical consistency.

Author Contributions: Conceptualization, H.W. and Y.L.; methodology, H.W.; software, H.W.; validation, H.W. and J.F.; formal analysis, H.W.; investigation, H.W.; data curation, H.W.; writing—original draft preparation, H.W.; writing—review and editing, J.F. and Y.L.; visualization, H.W.; supervision, Y.L.; project administration, Y.L.; funding acquisition, Y.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Research Fund of National Key Laboratory of Aerospace Physics in Fluids, grant number 2025-APF-KFZD-30; the Central Government Guides Local Projects of China, grant number 2025ZYDF105; and the International Scientific and Technological Cooperation Projects of Sichuan Province, grant number 2025YFHZ0148.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The 2D transonic RAE2822 airfoil dataset used in this study is publicly available in the Aero-Nef GitLab repository at <https://gitlab.isae-supaero.fr/gi.catalani/aero-nepf> (accessed on 25 August 2026) and is described in Ref. [27]. The source code developed for the present study is not publicly available.

Conflicts of Interest: The authors declare no conflicts of interest.

References

- Harten, A. High resolution schemes for hyperbolic conservation laws. *J. Comput. Phys.* **1983**, *49*, 357–393. [https://doi.org/10.1016/0021-9991\(83\)90136-5](https://doi.org/10.1016/0021-9991(83)90136-5).
- Spalart, P.R.; Allmaras, S.R. A one-equation turbulence model for aerodynamic flows. In *Proceedings of the 30th Aerospace Sciences Meeting and Exhibit*; AIAA: Reston, VA, USA, 1992; AIAA-92-0439. <https://doi.org/10.2514/6.1992-439>.
- Thurey, N.; Weissenow, K.; Prantl, L.; Hu, X. Deep learning methods for Reynolds-averaged Navier–Stokes simulations of airfoil flows. *AIAA J.* **2020**, *58*, 25–36. <https://doi.org/10.2514/1.J058291>.
- Fukami, K.; Fukagata, K.; Taira, K. Super-resolution reconstruction of turbulent flows with machine learning. *J. Fluid Mech.* **2019**, *870*, 106–120. <https://doi.org/10.1017/jfm.2019.238>.
- Deng, Z.; He, C.; Liu, Y.; Kim, K.C. Super-resolution reconstruction of turbulent velocity fields using a generative adversarial network-based artificial intelligence framework. *Phys. Fluids* **2019**, *31*, 125111. <https://doi.org/10.1063/1.5127031>.
- Fukami, K.; Fukagata, K.; Taira, K. Machine-learning-based spatio-temporal super resolution reconstruction of turbulent flows. *J. Fluid Mech.* **2021**, *909*, A9. <https://doi.org/10.1017/jfm.2020.948>.
- Kim, H.; Kim, J.; Won, S.; Lee, C. Unsupervised deep learning for super-resolution reconstruction of turbulence. *J. Fluid Mech.* **2021**, *910*, A29. <https://doi.org/10.1017/jfm.2020.1028>.
- Dong, C.; Loy, C.C.; He, K.; Tang, X. Image super-resolution using deep convolutional networks. *IEEE Trans. Pattern Anal. Mach. Intell.* **2016**, *38*, 295–307. <https://doi.org/10.1109/TPAMI.2015.2439281>.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; Lee, K.M. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*; IEEE: New York, NY, USA, 2017; pp. 1132–1140.
- Zhang, Y.; Tian, Y.; Kong, Y.; Zhong, B.; Fu, Y. Residual dense network for image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2018; pp. 2472–2481.
- Gao, H.; Sun, L.; Wang, J.X. PhyGeoNet: Physics-informed geometry-adaptive convolutional neural networks for solving parameterized steady-state PDEs on irregular domain. *J. Comput. Phys.* **2021**, *428*, 110079. <https://doi.org/10.1016/j.jcp.2020.110079>.
- Sitzmann, V.; Martel, J.N.P.; Bergman, A.W.; Lindell, D.B.; Wetzstein, G. Implicit neural representations with periodic activation functions. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 7462–7473.
- Mildenhall, B.; Srinivasan, P.P.; Tancik, M.; Barron, J.T.; Ramamoorthi, R.; Ng, R. NeRF: Representing scenes as neural radiance fields for view synthesis. In *Computer Vision—ECCV 2020*; Springer: Cham, Switzerland, 2020; pp. 405–421.
- Park, J.J.; Florence, P.; Straub, J.; Newcombe, R.; Lovegrove, S. DeepSDF: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2019; pp. 165–174.
- Ha, D.; Dai, A.; Le, Q.V. HyperNetworks. In *Proceedings of the International Conference on Learning Representations*, Toulon, France, 24–26 April 2017.
- Chen, Y.; Liu, S.; Wang, X. Learning continuous image representation with local implicit image function. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2021; pp. 8628–8638.
- Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **2019**, *378*, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
- Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. <https://doi.org/10.1038/s42254-021-00314-5>.
- Pfaff, T.; Fortunato, M.; Sanchez-Gonzalez, A.; Battaglia, P.W. Learning mesh-based simulation with graph networks. In *Proceedings of the International Conference on Learning Representations*, Virtually, 3–7 May 2021.
- Tancik, M.; Srinivasan, P.P.; Mildenhall, B.; Fridovich-Keil, S.; Raghavan, N.; Singhal, U.; Ramamoorthi, R.; Barron, J.; Ng, R. Fourier features let networks learn high frequency functions in low dimensional domains. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 7537–7547.
- Li, Z.; Kovachki, N.; Azizzadenesheli, K.; Liu, B.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Fourier neural operator for parametric partial differential equations. In *Proceedings of the International Conference on Learning Representations*, Virtually, 3–7 May 2021.
- Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; Karniadakis, G.E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* **2021**, *3*, 218–229. <https://doi.org/10.1038/s42256-021-00302-5>.
- Kovachki, N.; Li, Z.; Liu, B.; Azizzadenesheli, K.; Bhattacharya, K.; Stuart, A.; Anandkumar, A. Neural operator: Learning maps between function spaces with applications to PDEs. *J. Mach. Learn. Res.* **2023**, *24*, 1–97.
- Roe, P.L. Approximate Riemann solvers, parameter vectors, and difference schemes. *J. Comput. Phys.* **1981**, *43*, 357–372. [https://doi.org/10.1016/0021-9991\(81\)90128-5](https://doi.org/10.1016/0021-9991(81)90128-5).
- Jiang, G.S.; Shu, C.W. Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* **1996**, *126*, 202–228. <https://doi.org/10.1006/jcph.1996.0130>.

26. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2016; pp. 770–778.
27. Catalani, G.; Agarwal, S.; Bertrand, X.; Tost, F.; Bauerheim, M.; Morlier, J. Neural fields for rapid aircraft aerodynamics simulations. *Sci. Rep.* **2024**, *14*, 25496. <https://doi.org/10.1038/s41598-024-76983-w>.
28. Cook, P.H.; McDonald, M.A.; Firmin, M.C.P. *Aerofoil RAE 2822: Pressure Distributions, and Boundary Layer and Wake Measurements*; AGARD Report AR 138, Experimental Data Base for Computer Program Assessment; Royal Aircraft Establishment: Hampshire, UK, 1979.
29. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **2019**, *32*, 8024–8035.
30. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations*, San Diego, CA, USA, 7–9 May 2015.
31. Loshchilov, I.; Hutter, F. SGDR: Stochastic gradient descent with warm restarts. In *Proceedings of the International Conference on Learning Representations*, Toulon, France, 24–26 April 2017.
32. Sanchez-Gonzalez, A.; Godwin, J.; Pfaff, T.; Ying, R.; Leskovec, J.; Battaglia, P.W. Learning to simulate complex physics with graph networks. In *Proceedings of the 37th International Conference on Machine Learning*, Virtually, 12–18 July 2020; pp. 8459–8468.
33. Gao, H.; Ji, S. Graph U-Nets. In *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, CA USA, 9–15 June 2019; pp. 2083–2092.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.