A High-Precision Airfoil Pressure Distribution Prediction Model Based on Hybrid CNN and Chebyshev-Kolmogorov-Arnold Network

CHEN Yaohong[1], *ZHANG Luchi*[2], *DENG Yiju*[3], *YU Yanze*[3], *LI Xiang*[4], *JIAO Renshan*[1*]

1. Harbin Branch, AVIC Aerodynamics Research Institute, Harbin 150001, P. R. China; 2. China Gas Turbine Establishment, Chengdu 610599, P. R. China; 3. Graduate School, Chinese Aeronautical Establishment, Yangzhou 225111, P. R. China; 4. Shenyang Branch, AVIC Aerodynamics Research Institute, Shenyang 110034, P. R. China

(Received 30 May 2026; revised 28 July 2026; accepted 13 August 2026)

Abstract: To address the high computational costs of traditional computational fluid dynamics (CFD) in aerodynamic optimization and the limitations of conventional deep learning models, such as slow convergence and insufficient accuracy in capturing complex flow physics, a novel framework integrating convolutional neural network (CNN) and Chebyshev-Kolmogorov-Arnold network (Cheby-KAN) is proposed for 2D airfoil pressure distribution prediction. The framework employs a CNN-based geometry encoder to automatically extract high-dimensional features from signed distance field (SDF) and map them into low-dimensional latent vectors, thereby bypassing traditional manual parameterization. In the regression module, shifted Chebyshev polynomials are introduced as basis functions within the KAN architecture (Cheby-KAN). Leveraging their orthogonality and superior approximation properties, Cheby-KAN effectively mitigates the Runge phenomenon and enhances numerical stability when resolving high-gradient pressure patterns. Experimental results on a dataset comprising 89 million data points demonstrate that Cheby-KAN achieves a 60% reduction in training time compared to conventional multi-layer perceptron (MLP), with a 92% decrease in training mean square error (MSE). Furthermore, the model exhibits exceptional generalization in "zero-shot" tests on unseen airfoils, maintaining a mean absolute error (MAE) of 2.1×10^{-3}, which underscores its robustness as a high-fidelity surrogate model for advanced aerodynamic design.

Key words: airfoil pressure distribution; deep learning; convolutional neural network (CNN); Chebyshev polynomials; Kolmogorov-Arnold network (KAN); aerodynamic prediction

CLC number: V211.3 **Document code:** A **Article ID:** 1005-1120(2026)04-0592-22

0 Introduction

Airfoil design is fundamental to modern aerospace, wind energy, and gas turbine industries[1]. High-fidelity characterization of the surrounding flow field—specifically pressure distributions—serves as the prerequisite for optimizing aerodynamic performance[2-3]. Such data are indispensable for quantifying critical parameters like lift and drag, which directly govern mechanical efficiency, structural integrity, and operational longevity[4].

Traditional acquisition of the flow field surrounding an airfoil relies on solving the Navier-Stokes (N-S) equations, a set of partial differential equations that describe fluid motion. Due to their high nonlinearity and complexity, solving these equations typically necessitates significant computational resources and time[5]. Consequently, computational fluid dynamics (CFD) techniques are employed. Although CFD techniques enable the nu-

*Corresponding author, E-mail address: marlboroiron@163.com.

How to cite this article: CHEN Yaohong, ZHANG Luchi, DENG Yiju, et al. A high-precision airfoil pressure distribution prediction model based on hybrid CNN and Chebyshev-Kolmogorov-Arnold network[J]. Transactions of Nanjing University of Aeronautics and Astronautics, 2026, 43(4): 592-613.

http://dx.doi.org/10.16356/j.1005-1120.2026.04.008

merical solution of N-S equations on discrete meshes, they are characterized by high computational costs and excessive time consumption. Furthermore, during large-scale parametric sweeps or when resolving complex flow phenomena, the computational cost of CFD increases exponentially[6].

Breakthroughs in data science and deep learning are driving a paradigm shift in flow field modeling, transitioning from purely physics-driven to data-driven surrogates. These emerging methodologies show great promise in overcoming the limitations of traditional methods[7]. Deep learning models can extract underlying patterns from large-scale simulation datasets, thereby enabling rapid flow field prediction[8]. This transition not only boosts throughput but also offers a novel lens to resolve intricate flow features that remain computationally intractable for conventional techniques[9]. For instance, Obiols-Sales et al.[10] developed the CFDNet framework, which integrated convolutional neural network (CNN) to efficiently solve the Reynolds-averaged Navier-Stokes (RANS) equations. This approach achieves a speedup in flow field prediction of two orders of magnitude compared to traditional methods. Mi et al.[11] proposed an enhanced U-Net-LSTM hybrid neural network that utilized the signed distance function (SDF) and composite imagery to augment low-dimensional geometric and operational condition information. This approach enables the high-efficiency modeling of both steady and unsteady flow fields. Experimental results demonstrate that the network achieves prediction errors of less than 1.98% and 2.56% under steady and unsteady conditions, respectively. Thuerey et al.[12] systematically evaluated the predictive capabilities of various neural networks for pressure and velocity fields based on an enhanced U-Net architecture. Their research revealed the impact of training dataset scale and model parameters on predictive accuracy, achieving a mean relative error of less than 3% across diverse unseen airfoil configurations. Tompson et al.[13] proposed an innovative approach that combines local convolutional operations with global downsampling and upsampling structures. This method leverages CNN to capture both long-range physical phenomena (e.g., pressure gradients) and local flow details. Compared to traditional U-Net models, this approach significantly enhances the modeling of long-range effects, thereby better capturing long-distance pressure propagation—A critical limitation of local convolutions in standard U-Net architectures. Furthermore, numerous researchers[14-16] have developed surrogate models for pressure distribution prediction based on various machine learning methods, applying them to establish efficient aerodynamic design frameworks. Notwithstanding these advancements, conventional deep learning architectures for flow field prediction are still hindered by several critical bottlenecks. In terms of fidelity, a persistent discrepancy exists between neural-network-inferred fields and high-fidelity CFD benchmarks. This is particularly evident in the leading and trailing edge regions, where infinitesimal prediction errors can be amplified by differential operators, leading to substantial biases in gradient-based aerodynamic coefficients[17]. Regarding robustness, existing models often struggle to maintain both generalization and accuracy when trained on sparse datasets. Furthermore, computational efficiency remains a major concern, as the vast parameterization of standard neural networks entails prohibitive training costs.

To address these issues, we propose a novel 2D airfoil flow field prediction framework, termed CNN-Cheby-KAN, which integrates CNN with Chebyshev-Kolmogorov-Arnold networks (Cheby-KAN). The CNN component eliminates the reliance on hand-crafted geometric features in traditional parameterization techniques, such as class-shape transformation (CST) and non-uniform rational B-splines (NURBS). KAN significantly reduces training complexity and parameters by representing complex multivariate functions as combinations of univariate functions[18]. Additionally, the integration of Chebyshev polynomials effectively mitigates numerical stability issues and suppresses the Runge phenomenon—The severe oscillations at interval edges typically associated with high-order polynomial interpolation. Experimental results demonstrate that, compared to multi-layer perceptron (MLP)-based methods, the proposed model exhibits superior predictive performance and generalization capabilities for pressure distributions around airfoils with significantly reduced computational time.

The remainder of this paper is organized as follows: Section 1 elucidates the fundamental principles underlying this study. Section 2 outlines the methodology for dataset construction and discusses its validity. In Section 3, a comprehensive prediction model is developed using the airfoil dataset, followed by a comparative analysis of different models to highlight the advantages of the proposed CNN-Cheby-KAN architecture. This section also explicates the role of Chebyshev polynomials in enhancing model performance and evaluates the framework's generalization capabilities. Section 4 gives the experimental results. Finally, Section 5 concludes the study and offers suggestions for future research.

1 Preliminaries

1.1 Fundamental principles of CNNs

CNN is a class of deep learning architectures specifically engineered to process grid-like data structures, such as images or discretized flow field meshes. Their design is predicated on core principles including local connectivity, weight sharing, and spatial pooling, which collectively facilitate the efficient extraction of hierarchical features from raw input data. The complete architecture is illustrated in Fig.1, and ReLU represents the rectified linear unit.

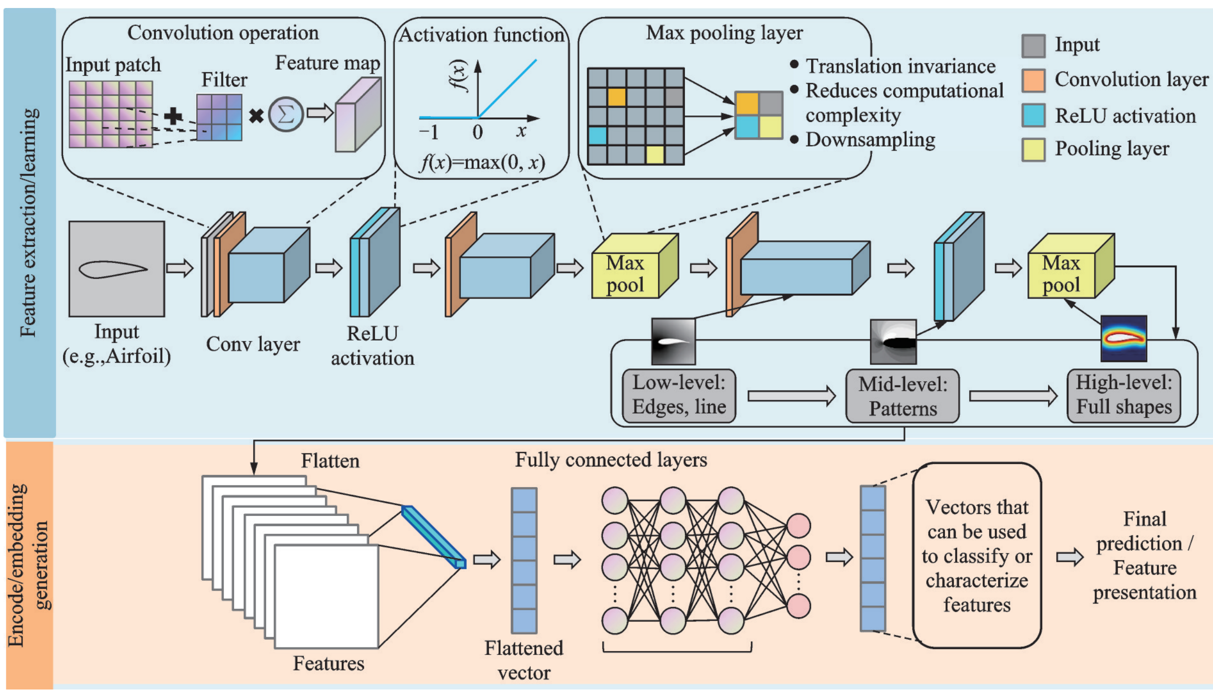

Fig.1 Schematic diagram of the complete CNN architecture

1.1.1 Convolutional layers and feature extraction

The convolutional layer performs sliding window operations on the input signal using a set of learnable filters (also referred to as kernels). This architecture not only drastically reduces the number of parameters but also enables the precise capture of local features with translation invariance. For a 2D input image $\boldsymbol{X}$ with the height H, width W, and C channels, the value of the output feature map $\boldsymbol{Y}$ at position (i,j) is computed using a convolutional kernel $\boldsymbol{K}$ as follows[19-20]

$$Y_k(i,j)=\sigma\left(\sum_{c=1}^{C}\sum_{m=1}^{K_\mathrm{h}}\sum_{n=1}^{K_\mathrm{w}}\boldsymbol{X}(i\cdot s+m,j\cdot s+n,c)\cdot W_k(m,n,c)+b_k\right) \tag{1}$$

where $Y_k(i,j)$ represents the value of the kth output feature map at spatial location (i,j), K_h and K_w represent the height and width of the convolutional kernel, s is the stride of the convolutional kernel, m and n are the indices within the kernel, W_k is the weight of the kernel at that position, b_k is the bias term, and σ represents the non-linear activation function. By applying multiple convolutional kernels in parallel, the network can simultaneously extract a diverse array of features, such as gradients or edges within the flow field. Mathematically, the convolution operation can be viewed as a learnable discrete differential operator, which is particularly effective for identifying high-gradient regions near the airfoil surface.

1.1.2 Pooling operations and dimensionality reduction

Following the convolutional operations, pooling layers are typically employed to reduce the spatial dimensions of the feature maps. By selecting either the maximum value (max pooling) or the average value (average pooling) within a specified window[21], this operation effectively lowers the data sampling rate while preserving the most prominent features. This mechanism enhances the model's robustness against geometric distortions, ensuring more stable performance when processing complex airfoil deformations or mesh warping. Beyond dimensionality reduction, pooling operations facilitate the progressive expansion of the receptive field, enabling the network to aggregate local aerodynamic details into a global understanding of the flow field topology.

1.1.3 Non-linear activation and deep representation

To model the intricate non-linear mappings inherent in fluid flow, CNN incorporates non-linear activation functions, the most prevalent of which is ReLU[22], shown as

$$\mathrm{ReLU}(x)=\max(0,x) \tag{2}$$

By stacking multiple convolutional, pooling, and activation layers, the network can progressively abstract high-level flow patterns from low-level geometric contours. This hierarchical feature extraction capability provides CNN with a distinct advantage in establishing low-dimensional embeddings and feature representations of the external flow fields around airfoils. The resulting compact feature representation serves as a robust input for the subsequent regression module, ensuring that the critical geometric constraints are preserved throughout the mapping process.

1.2 Fundamental principles of KAN

1.2.1 Mathematical background: Kolmogorov-Arnold representation theorem

KAN represents a novel neural network architecture inspired by the Kolmogorov-Arnold representation theorem. This theorem asserts that any multivariate continuous function $f:[0,1]^n\to\mathbf{R}$ defined on a bounded domain can be represented as a finite composition of continuous univariate functions and addition operations. Specifically, the theorem guarantees that the function f can be expressed in the following form[23]

$$f(x_1,x_2,\cdots,x_n)=\sum_{q=1}^{2n+1}\Phi_q\left(\sum_{p=1}^{n}\phi_{q,p}(x_p)\right) \tag{3}$$

where each $\phi_{q,p}$ and Φ_q are continuous univariate functions. This result indicates that the structure of any multivariate continuous function can be fundamentally decomposed into the addition and composition of univariate functions. The complete architecture of KAN is illustrated in Fig.2.

1.2.2 The fundamental architecture of KAN

Although the original Kolmogorov-Arnold representation theorem is mathematically complete, its direct implementation in modern deep learning is hindered by the potential non-smooth or fractal nature of the univariate functions, coupled with its restricted two-layer structure. Inspired by the work in Ref.[18], KAN bypasses these limitations by relaxing the constraints on the number of intermediate nodes and allowing the construction of arbitrary network depths.

A KAN layer with n_{in} input nodes and n_{out} output nodes is defined as a matrix of 1D learnable functions, shown as

$$\Phi=\{\phi_{j,i}\}\quad i=1,2,\cdots,n_{\text{in}};j=1,2,\cdots,n_{\text{out}} \tag{4}$$

For a deep KAN architecture, forward propagation is achieved through the composition of multiple KAN layers. The value of the jth node in the $(l+1)$th layer, denoted as $x_j^{(l+1)}$, is computed as

$$x_j^{(l+1)}=\sum_{i=1}^{n_l}\phi_{l,j,i}(x_i^{(l)}) \tag{5}$$

In stark contrast to MLP, which apply fixed activation functions at the nodes, KAN positions

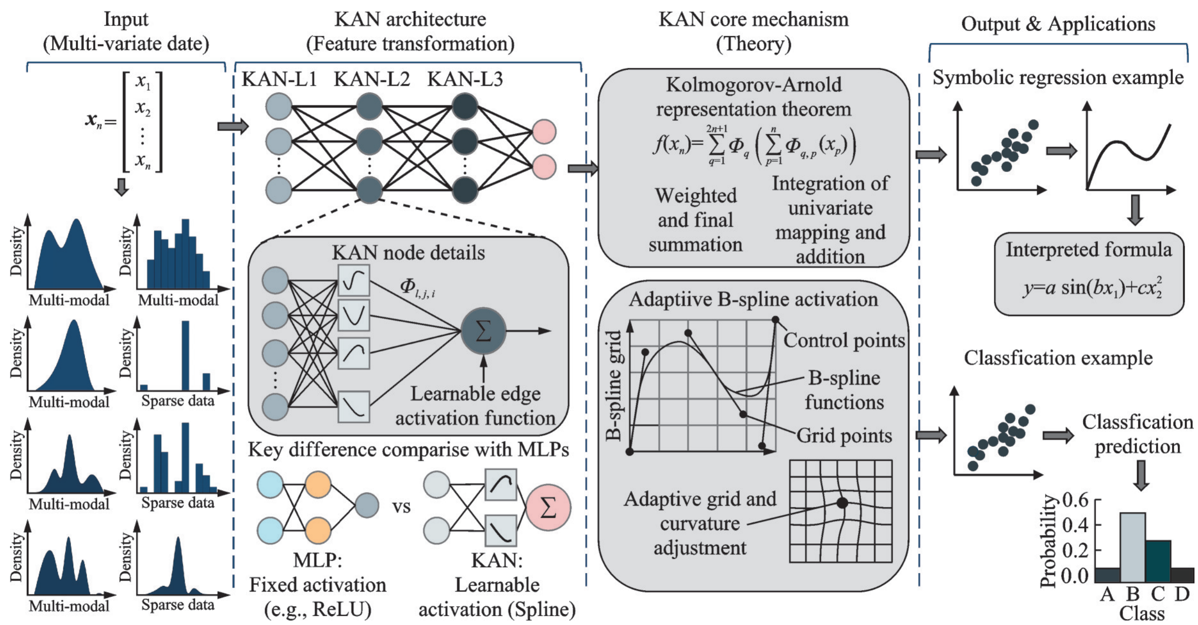

Fig.2 Architectural framework and operational mechanism of KAN

learnable univariate activation functions on the "edges" (links). The nodes themselves perform only simple summation operations, without any nonlinear transformations.

1.2.3 Function parameterization and implementation

Mathematically, each learnable activation function ϕ is constructed as a combination of a base function and a differentiable spline function, shown as

$$\phi(x)=w\cdot(b(x)+\mathrm{spline}(x)) \tag{6}$$

where the base function $b(x)$ is typically the Sigmoid linear unit (SiLU) function, shown as

$$\mathrm{SiLU}(x)=\frac{x}{1+\mathrm{e}^{-x}} \tag{7}$$

The spline component, $\mathrm{spline}(x)$, is represented as a linear combination of B-splines basis functions, shown as

$$\mathrm{spline}(x)=\sum_{i=0}^{N}c_i B_i(x) \tag{8}$$

where $B_i(x)$ denotes the ith B-spline basis function, and c_i the trainable coefficients. The SiLU function provides a fundamental global non-linearity, while the B-splines offer flexible and smooth local representation capabilities, enabling the model to capture complex, non-uniform patterns. However, B-splines can encounter challenges regarding numerical stability and grid-dependency when scaling to high-order approximations, motivating the exploration of alternative basis functions.

1.2.4 Performance advantages and interpretability

Existing studies[24-26] have demonstrated that KAN exhibit superior accuracy and parameter efficiency compared to traditional MLP in tasks such as data fitting and solving partial differential equations (PDEs). For instance, in PDE-solving benchmarks, a 2-layer KAN with a width of 10 achieves lower error (10^{-7} vs. 10^{-5}) and significantly higher parameter efficiency (10^2 vs. 10^5 parameters) compared to a 4-layer MLP with a width of 100, as summarized in Table 1 and Fig.3.

Table 1 Performance and parameter efficiency comparison between KAN and MLP

Model architecture	Layer	Width	Parameter	Error
KAN(Proposed, base)	2	10	$\approx 2.5\times 10^2$	10^{-7}
KAN(Proposed, small)	2	5	$\approx 10^2$	10^{-5}
MLP(Baseline)	4	100	$\approx 10^5$	10^{-5}

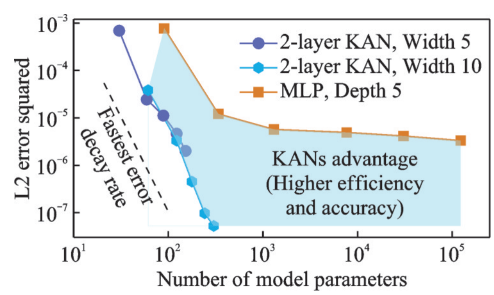

Fig.3 Scaling law comparison between KAN and MLP L2 error decay relative to the number of model parameters in PDE solving tasks

Furthermore, as the number of parameters increases, the test error of KAN decays significantly faster than that of MLPs, following a steeper scaling law. Beyond these performance gains, KAN offers exceptional interpretability: By applying symbolic regression and pruning to the edge functions, researchers can transform a "black-box" neural network into explicit mathematical formulas. This transparency holds immense potential for Scientific Discovery, enabling the extraction of underlying physical laws from data-driven models.

1.3 Improvement of KAN based on shifted Chebyshev polynomials

1.3.1 Motivation: Computational overhead and resource bottlenecks

The original B-spline-based KAN models face severe challenges when processing large-scale scientific computing datasets. In the airfoil external flow field prediction task addressed in this study, the flow field for each specific operating condition contains 262 144 spatial points, while each airfoil is characterized by 16 feature vectors. Given such a massive data scale, the unique architecture of KAN, which places activation functions on the "edges", imposes a heavy computational burden. Specifically, the intermediate activation values and gradient information for each layer must maintain a computational graph for the spline basis functions on every single edge. This leads to a memory footprint that scales exponentially with the data volume. Furthermore, B-splines rely on local grid indexing and addressing; such fragmented operations fail to fully leverage the tensor cores of modern GPUs, resulting in training durations that significantly exceed practical expectations.

1.3.2 Introduction of shifted Chebyshev polynomials

To surmount the aforementioned deficiencies, we replace B-splines with shifted Chebyshev polynomials, which possess superior algebraic properties. Given that physical parameters in aerodynamic simulations are typically normalized to non-negative intervals, we employ the shifted polynomials $T_n^*(x)$ defined on $x\in[0,1]$, rather than the classical definition on $[-1,1]$.

For an nth order polynomial, the shifted Chebyshev polynomials can be defined via the following recurrence relation[27], shown as

$$\begin{cases} T_0^*(x)=1 \\ T_1^*(x)=2x-1 \\ T_{n+1}^*(x)=2(2x-1)T_n^*(x)-T_{n-1}^*(x) \end{cases} \tag{9}$$

By leveraging this recursive structure, the computation of $\phi(x)$ is transformed into highly efficient, large-scale matrix multiplications. This enables a high degree of computational vectorization, completely bypassing the cumbersome local grid update procedures inherent in the original KAN architecture[28].

1.3.3 Mathematical expression and parameterization

In the Cheby-KAN architecture, the nonlinear mapping $\phi_{j,i}(x)$ connecting node i in layer l to node j in layer $l+1$ is expressed as

$$\phi_{j,i}(x)=\sum_{n=0}^{N} c_{j,i,n}\cdot T_n^*(x) \tag{10}$$

where $N=5$ is the predefined degree (corresponding to six basis functions), and $c_{j,i,n}$ the learnable weight coefficients. Based on preliminary sensitivity analysis, $N=5$ provides an optimal balance between functional expressivity and computational stability. We define the basis function expansion vector $\boldsymbol{\Gamma}(x)$ and the learnable parameter tensor $\boldsymbol{C}$ for each layer as follows

$$\boldsymbol{\Gamma}(x)=[T_0^*(x),T_1^*(x),\cdots,T_N^*(x)]^{\mathrm{T}}\in\mathbf{R}^{(N+1)\times 1} \tag{11}$$

$$\boldsymbol{C}\in\mathbf{R}^{n_{\mathrm{out}}\times n_{\mathrm{in}}\times(N+1)} \tag{12}$$

For the input vector $\boldsymbol{x}^{(l)}$ of the lth layer, the forward propagation calculation for the output node j is simplified to

$$x_j^{(l+1)}=\sum_{i=1}^{n_{\mathrm{in}}}(c_{j,i}\cdot\boldsymbol{\Gamma}(x_i^{(l)})) \tag{13}$$

1.3.4 Stability and numerical advantages

The adoption of Chebyshev polynomials is motivated not only by computational acceleration but also by their profound mathematical advantages. Unlike traditional monomial power series, the roots of Chebyshev polynomials are clustered toward the interval boundaries[29]. This distribution effectively suppresses the severe oscillations typically observed in high-order interpolation near the edges, a phenomenon known as the Runge phenomenon, thereby ensuring the stability of cross-layer function approximation. Furthermore, due to the equioscillation property, Chebyshev polynomials yield an exceptionally uniform distribution of approximation errors within the interval [0, 1], shown as

$$\max_{x\in[0,1]}|f(x)-P_N(x)|\rightarrow \text{minimized} \tag{14}$$

Based on the orthogonality with respect to the weight function $w(x)=\dfrac{1}{\sqrt{x-x^2}}$, we have

$$\int_0^1 T_m^*(x)T_n^*(x)\frac{1}{\sqrt{x-x^2}}\mathrm{d}x=\begin{cases}0 & m\neq n\\ \dfrac{\pi}{2} & m=n>0\\ \pi & m=n=0\end{cases} \tag{15}$$

This orthogonality significantly reduces the gradient coupling between coefficients of different orders, enhancing numerical stability during backpropagation. This characteristic is paramount for data modeling in CFD, where sensitivity to precision is extremely high. The architectural integration and the resulting numerical robustness are summarized in Fig.4.

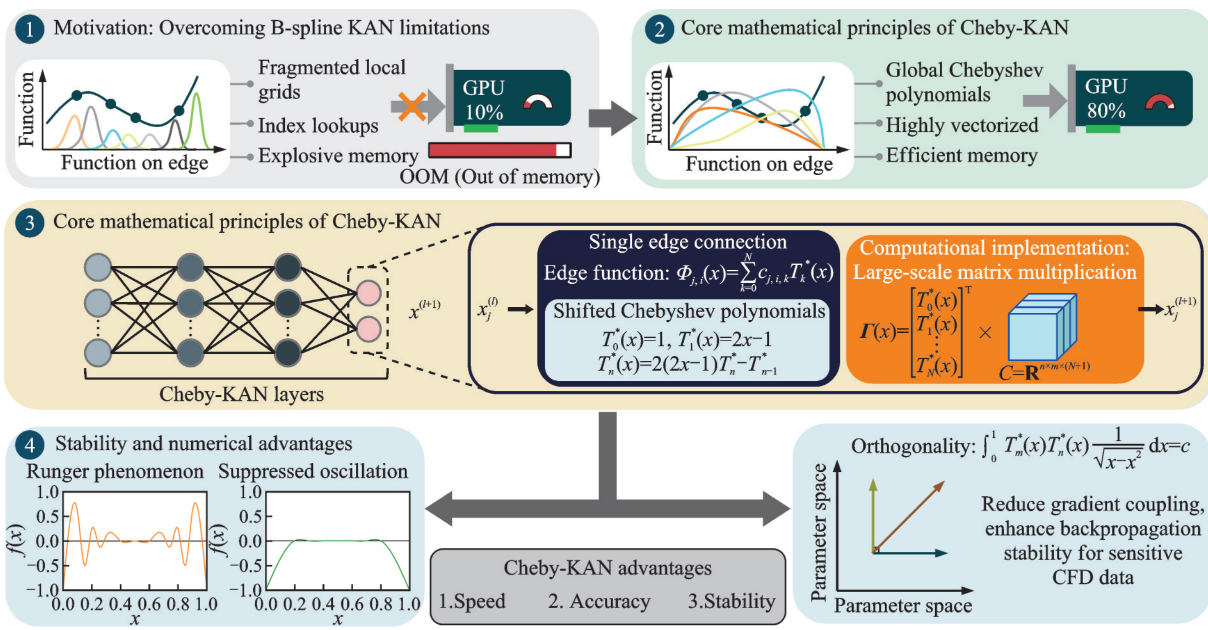

Fig.4 Architecture and mathematical principles of Cheby-KAN

2 Data Preparation

2.1 Airfoil data preparation

Conventional methodologies for airfoil geometry construction primarily utilize NACA series parameterization, CST, or B-spline/NURBS curve representations. However, the reliance on fixed basis functions in CST or NURBS often necessitates a compromise between global smoothness and local fidelity, which is particularly problematic when resolving high-curvature regions such as the leading edge.

In contrast, CNN possesses an intrinsic advantage in capturing these locally translation-invariant features. By leveraging a grid-based representation, CNN can directly extract high-dimensional latent features from the airfoil contour, effectively bypassing the information loss inherent in traditional dimensionality reduction and manual parameterization.

2.1.1 Dataset standardization and preprocessing

The UIUC Airfoil Data Site contains approximately 1 600 airfoil geometries, provided as sets of discrete coordinate points[30]. However, the raw data points often exhibit non-uniform density and potential numerical noise, making them unsuitable for direct input into a neural network. To address this, we established a standardized preprocessing pipeline.

First, B-spline curves are employed to fit and smooth the original discrete coordinates. This step not only filters out infinitesimal noise but also leverages the recursive properties of B-splines to perform uniform high-density resampling (interpolation, 350 curvature-adaptive points per airfoil) along the airfoil contour. This ensures that all airfoils possess identical geometric resolution and spatial continuity.

Subsequently, to eliminate physical scale discrepancies and prevent absolute dimensions from biasing the model training, a geometric normalization strategy is applied. The core approach utilizes the chord length (c) as a unified reference scale to transform the coordinate data into non-dimensional relative quantities. The normalization equations are as follows

$$x_{\text{norm}} = \frac{x - x_{\min}}{c} \tag{16}$$

$$y_{\text{norm}} = \frac{y}{c} \tag{17}$$

where the normalized chord length is defined as $c = x_{\max} - x_{\min}$. Here, $x_{\max}$ and $x_{\min}$ represent the maximum and minimum longitudinal coordinates of the raw airfoil, respectively. Specifically, the inclusion of $x_{\min}$ ensures that the leading edge of every airfoil is strictly aligned at $x_{\text{norm}} = 0$. Through this procedure, critical features such as relative thickness and camber are preserved, while absolute size variations are eliminated. Furthermore, the airfoils are rotated to ensure the chord line aligns with the horizontal axis, minimizing geometric bias.

2.1.2 Signed distance field representation

Upon completing geometric preprocessing, the two-dimensional contour points should be transformed into a grid format compatible with neural networks. To fully capture local geometric features and their spatial influence, this study adopts the SDF as the geometric input representation for the CNN[31-32].

First, a high-resolution computational grid $\boldsymbol{X}$-$\boldsymbol{Y}$ of 512 pixel×512 pixel is generated. This grid not only encompasses the normalized airfoil itself ($x_{\text{norm}} \in [0, 1]$) but also extends into the surrounding area to include a sufficient portion of the external flow field. The grid covers a domain of $x \in [-1.0, 2.0]$ and $y \in [-1.5, 1.5]$ to ensure the leading-edge stagnation point and wake region are included. This extended domain allows the CNN to perceive the potential impact of the airfoil geometry on the global spatial positions.

The value of each pixel (i, j) in the SDF image represents the signed physical distance from that grid point to the nearest point on the normalized airfoil contour. Its mathematical expression is given by

$$\text{SDF}_{ij} = \text{sign}((X_{ij}, Y_{ij})) \cdot \min_{p \in P_{\text{norm}}} \| (X_{ij}, Y_{ij}) - p \|_2 \tag{18}$$

where P_{norm} is the set of normalized airfoil contour points and $\| \cdot \|_2$ the Euclidean distance. The sign function sign(•) is defined such that the value is negative (−) when the grid point is located inside the airfoil, positive (+) when outside, and zero (0) exactly on the contour. The visual transition from the discrete coordinate points to the continuous signed distance representation is illustrated in Fig.5.

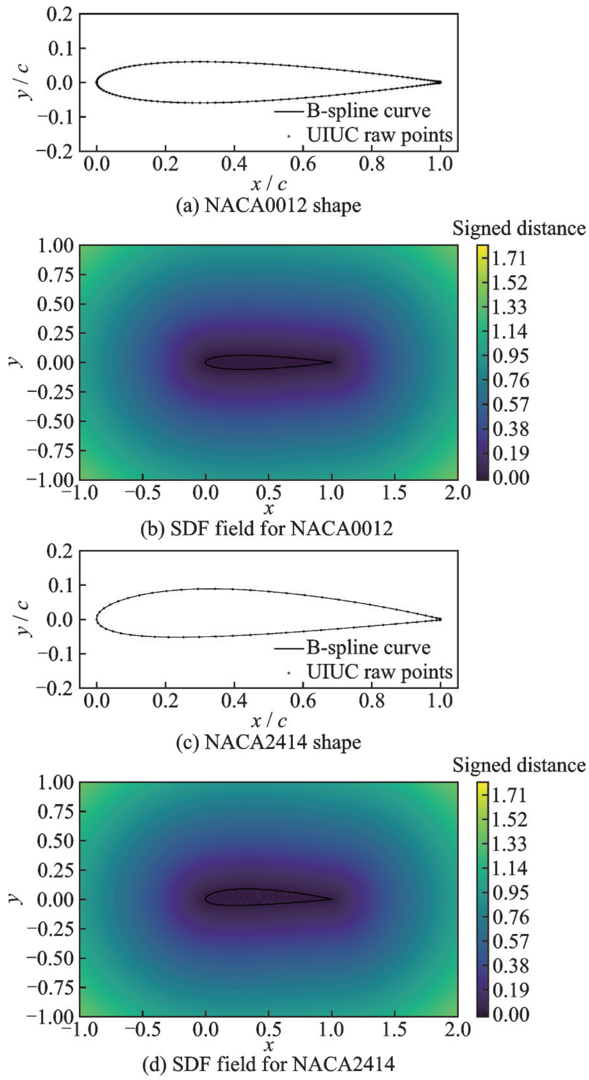

Fig.5 Comparison of airfoil contour and SDF representations

2.2 CNN-based automatic feature extraction of airfoil geometric features

Although the pixel-based SDF images generated through the standardization process precisely capture global geometric and spatial information, their high dimensionality (512 pixel × 512 pixel) poses a significant challenge for direct input into the subsequent complex flow field fitting network (Cheby-KAN). To achieve efficient end-to-end modeling for scientific computing, we utilize a CNN as a geometry encoder. This component facilitates active parameterization and dimensionality reduction, mapping the data from the high-dimensional image space to a low-dimensional latent space.

The primary structure and configuration of the CNN-based feature extractor are illustrated in Fig.6. By training the CNN to compress geometric information, the model identifies the most salient aerodynamic features, such as leading-edge curvature and maximum thickness distribution, and represents them as a compact feature vector. This latent representation serves as a robust input for the Cheby-KAN module, ensuring that the critical geometric constraints are preserved while significantly reducing the computational complexity of the subsequent regression task.

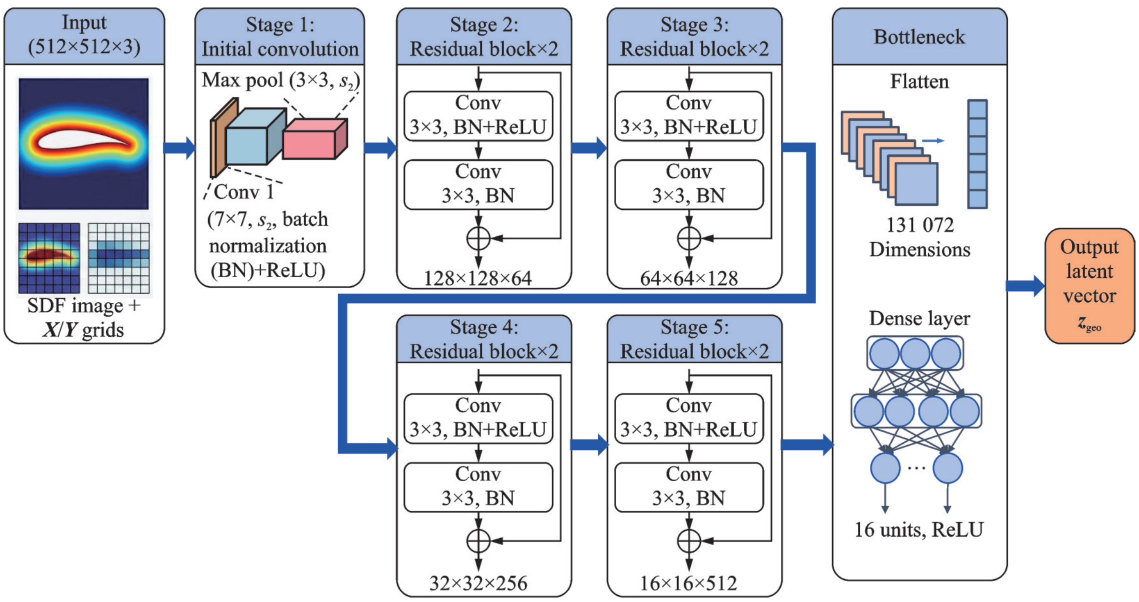

Fig.6 Main layers and core parameters of CNN parametric extractor

2.2.1 CNN architecture and training for geometric feature extraction

The input to the network is a tensor with dimensions (512 × 512 × 3), where the three channels carry spatial and geometric information. To preserve high-resolution features while mitigating the vanishing gradient problem associated with deep architectures, we employ a residual network (ResNet) framework. The specific hyperparameters of the convolutional network are summarized in Table 2.

Table 2 Related parameters of CNN airfoil parametric encoding

Stage	Layer/Module type	Kernel size	Channel	Output size	Activation
Input	SDF+Grid	—	3	512 × 512	—
Stage 1	Conv 1	7 × 7, stride=2	64	256 × 256	ReLU
—	Max pool	3 × 3, stride=2	64	128 × 128	—
Stage 2	Residual block×2	3 × 3	64	128 × 128	ReLU
Stage 3	Residual block×2	3 × 3, stride=2	128	64 × 64	ReLU
Stage 4	Residual block×2	3 × 3, stride=2	256	32 × 32	ReLU
Stage 5	Residual block×2	3 × 3, stride=2	512	16 × 16	ReLU
Neck	Global avg pool	—	512	1 × 1 × 512	—
Output	Dense layer (Flatten)	—	16	16	ReLU

The complete encoder consists of five main convolutional stages, totaling 18 layers with learnable parameters (including convolutional and fully connected layers). Each convolutional layer is followed by BN and a ReLU activation function to accelerate convergence and enhance non-linear expressivity.

The dense feature maps extracted through these five stages (with a final dimension of $16 \times 16 \times 512$) are processed via global average pooling (GAP) and subsequently flattened into a one-dimensional vector, which then passes through a fully connected (dense) layer with ReLU activation. The network ultimately outputs a $d_{\text{latent}} = 16$ feature vector $\boldsymbol{z}_{\text{geo}}$.

2.2.2 Self-supervised pre-training and generalization

The CNN geometry encoder is trained in a self-supervised manner by minimizing the mean squared error (MSE) between the original SDF image and the reconstructed image from a decoder. Fig.7 illustrates the training loss curves. Owing to the dense SDF representation and the residual architecture, the loss exhibits an exponential decay within the first 10 epochs, indicating that the network rapidly captures the dominant geometric features. Following a stepped learning rate decay, the loss stabilizes and converges after 60 epochs, achieving a final validation MSE of 10^{-5}. This confirms that the extracted 16-dimensional latent vector $\boldsymbol{z}_{\text{geo}}$ encapsulates sufficient information to accurately reconstruct critical geometric details, successfully achieving active parameterization and efficient dimensionality reduction.

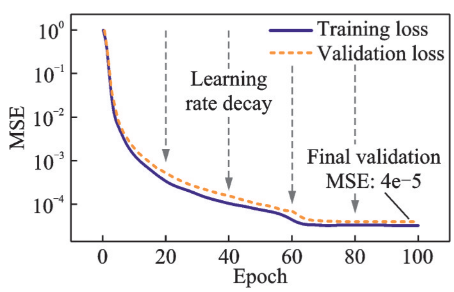

Fig.7 Training and validation loss curves of the CNN geometry encoder

Once the self-supervised pre-training of the CNN geometry encoder is completed, its network weights are strictly frozen. During the subsequent training phase of the flow field regressor (Cheby-KAN), only the parameters of the Cheby-KAN are updated, utilizing the frozen CNN solely as an offline feature extractor to generate the latent vector $\boldsymbol{z}_{\text{geo}}$.

To verify that the compressed 16-dimensional geometric latent space $\boldsymbol{z}_{\text{geo}}$ preserves critical geometric features without loss of high-fidelity details, a rigorous reconstruction error analysis is performed on several representative airfoils. Fig.8 illustrates the comparison between the ground-truth airfoil contours and the reconstructed zero-level sets extracted from the predicted SDF, covering standard conventional baselines (the symmetric e168 and asymmetric NACA4412), an extreme high-camber profile (S1223), and a reflexed camber profile featuring local reverse curvature (NACA M24). The reconstructed profiles exhibit exceptional agreement with the original geometries, successfully capturing these challenging high-gradient characteristics without any unphysical smoothing. Quantitatively, the maximum geometric error remains below 2.32×10^{-4}

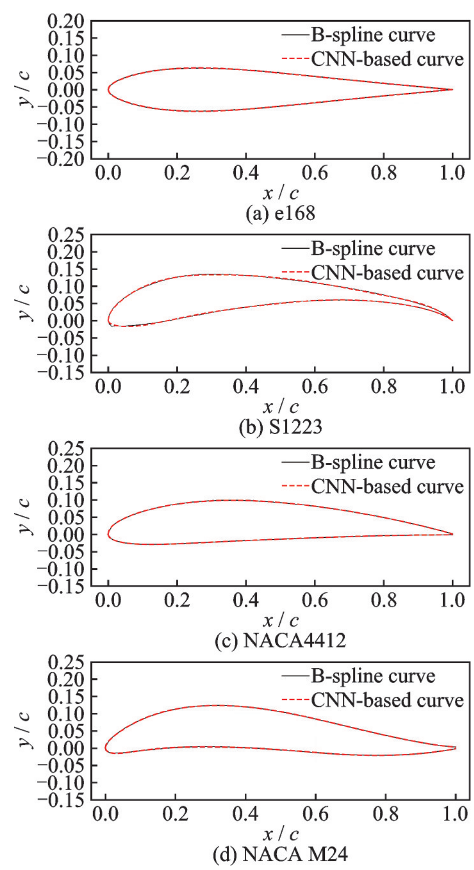

Fig.8 Original airfoil curves and the airfoil curves reconstructed using CNN

across all test cases. This demonstrably proves that the self-supervised CNN encoder possesses sufficient capacity to retain high-frequency geometric details, providing a mathematically accurate foundation for the subsequent Cheby-KAN regression.

2.3 Numerical simulation and dataset generation

In this study, several widely recognized airfoils, including the RAE2822, NACA0012, and e387, were selected to construct the dataset. These airfoils serve as established benchmarks in aerodynamic research[33-35], with their flow characteristics extensively documented in existing literature, making them ideal for validating our CFD setup and the proposed predictive model.

The characteristic chord length for all airfoils was standardized at 0.15 m. The far-field boundary conditions were prescribed with a static pressure $p_\infty = 101\,325$ Pa and a static temperature $T_\infty = 288.15$ K. Numerical simulations were conducted using the ANSYS Fluent solver to ensure robustness across diverse flow conditions. The shear stress transport (SST) k-ω turbulence model was employed, as it provides superior fidelity in capturing complex boundary layer development and potential flow separation on airfoil surfaces.

Prior to large-scale data generation, a rigorous grid independence study was performed on the NACA0012 airfoil. As illustrated in Fig.9, the variation in the lift coefficient C_l stabilized as the mesh density increased from 1.0×10^6 to 3.0×10^6 cells. To balance computational accuracy with efficiency, a final mesh resolution of approximately 2.0×10^6 units per airfoil was selected, ensuring sufficient resolution of flow field details, particularly in the critical leading and trailing edge regions.

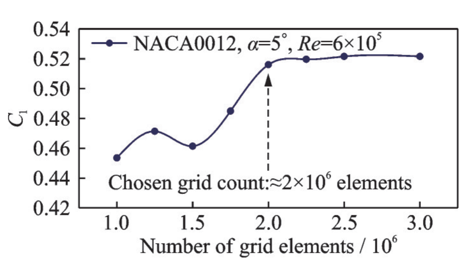

Fig.9 Grid independence study for NACA0012 airfoil

To further establish the credibility of the CFD dataset used for neural network training, the adopted numerical setup was validated using the benchmark NACA0012 airfoil at a Reynolds number of 5×10^5. Fig.10 compares the variation of the lift coefficient with angle of attack (AOA), α, obtained from the present CFD simulations against wind tunnel measurements[36]. Good agreement is observed throughout the pre-stall region, while only minor deviations appear at higher AOA, where flow separation becomes increasingly sensitive to turbulence-model assumptions. Since the lift coefficient is obtained by integrating the surface pressure distribution, the close agreement with experimental data also indirectly confirms the physical reliability of the generated pressure-field dataset. Therefore, the adopted CFD methodology provides a sufficiently accurate benchmark for constructing the neural network training database.

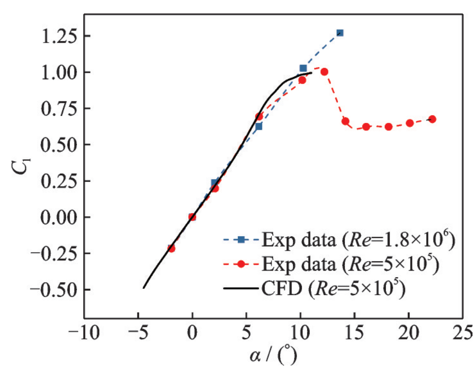

Fig.10 Comparison of lift coefficient characteristic curves between CFD simulation and experiment

For each airfoil, 80 distinct operating conditions were randomly sampled using Latin hypercube sampling (LHS) within a predefined velocity range (35 m/s to 75 m/s) and an AOA range (α: $-4.5°$ to $11.5°$). This approach ensures a space-filling and uniform distribution across the Reynolds number ($Re \in [4\times10^5, 8\times10^5]$) and α domains. For each airfoil, 80 operating conditions were generated using LHS, covering a velocity range of 35—75 m/s, an AOA range of $-4.5°$—$11.5°$, corresponding to a Reynolds number range of $Re = 4\times10^5$—8×10^5 (based on the unit chord length). Additionally, to expand the training space of the dataset and enhance the model's generalization ability, 10 airfoils were randomly selected from the UIUC airfoil database, with 10 operating conditions per airfoil. The relationship between the case-level samples and the spatial point-wise samples, as well as the dataset splitting strategy, is rigorously defined to ensure data independence and prevent data leakage.

The total dataset contains N_{case} distinct aerodynamic cases (3 airfoils × 80 conditions + 10 airfoils × 10 conditions). For each individual case, the continuous CFD flow field is projected onto a standardized structured sampling grid of 512 point × 512 point, yielding $N_{\text{grid}} = 262\,144$ spatial evaluation points per case. Consequently, the total number of point-wise training entries fed into the Cheby-KAN regression module is calculated as

$$N_{\text{total}} = N_{\text{case}} + N_{\text{grid}} = 89\,128\,960 \tag{19}$$

To eliminate any possibility of data leakage via spatial interpolation, the dataset splitting is executed strictly at the airfoil/operating-condition level (case-level), rather than randomly shuffling the 8.9×10^7 point-wise samples. Specifically, 90% of the cases (306 samples) are assigned to the training set, and the remaining 10% (34 samples) form the validation set. Under this strategy, all 262 144 spatial points belonging to a single CFD case are kept together in the same split, ensuring that the model is validated only on entirely unseen global flow regimes. Furthermore, completely unseen airfoil geometries are strictly withheld as an independent test set for zero-shot generalization and extrapolation evaluations in Section 4.3, guaranteeing absolute sample independence.

To adapt the discrete CFD results for neural network input, a standardized data processing pipeline was implemented. First, a square measurement domain was defined around each airfoil, covering the critical near-field flow regions. Subsequently, a 512 × 512 structured sampling grid, consistent with the network's input resolution, was generated within this domain to interpolate and resample the raw CFD flow variables (e.g., pressure and velocity).

To maintain the physical consistency of the pressure distribution across different conditions and eliminate absolute scale discrepancies, the extracted absolute pressure was converted into the non-dimensional pressure coefficient (C_{p}). The transformation is defined as

$$C_{\text{p}} = \frac{p - p_{\infty}}{\frac{1}{2}\rho_{\infty} v_{\infty}^2} \tag{20}$$

where p and p_{∞} denote the local static pressure and far-field static pressure, respectively; v_{∞} represents the free-stream velocity, and ρ_{∞} the far-field density determined by the ideal gas equation of state.

Finally, based on the high-density sampled field data, a standardized pressure distribution contour map for network input was generated, as illustrated in Fig.11. This standardized visual representation allows the CNN to effectively learn the spatial correlations and aerodynamic gradients across the entire computational domain.

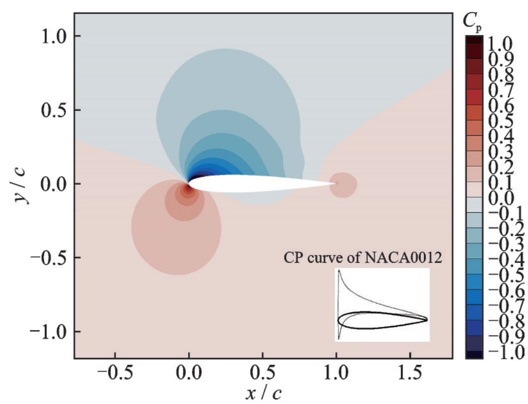

Fig.11 C_{p} distribution contour map for Cheby-KAN input (NACA0012)

3 Architecture of the Predictive Model

This section details the construction of the end-to-end neural network model, which integrates the geometric feature extractor (CNN encoder) with the physical flow field regressor (Cheby-KAN). Furthermore, the model's predictive performance and generalization capabilities are evaluated and analyzed using the CFD-generated datasets.

The input layer of the Cheby-KAN is designed to organically integrate key heterogeneous information influencing the flow field distribution, providing a comprehensive physical context for the network. The input layer receives a total of 20 parameters, mathematically represented as the fusion vector $\boldsymbol{V}_{\text{in}}$, shown as

$$\boldsymbol{V}_{\text{in}}=[\underbrace{x,y,}_{\text{Spatial position}}\underbrace{Re,\alpha,}_{\text{Flow conditions}}\underbrace{z_1,z_2,\cdots,z_{16}}_{\text{Geometric embedding } \boldsymbol{z}_{\text{geo}}}]^{\text{T}}\in\mathbf{R}^{20\times 1} \tag{21}$$

Specifically, the spatial coordinates (x,y) (two parameters) anchor the prediction targets to specific physical locations within the 512×512 standardized sampling grid. The flow condition parameters (Re,α) (two parameters), including the Reynolds number and the AOA, define the global flow regime within the physical domain. Finally, the 16-dimensional geometric feature vector $\boldsymbol{z}_{\text{geo}}$ (16 parameters) is automatically extracted by the CNN encoder from the dense $512\times512\times3$ (SDF + Grid) tensors. This latent representation encapsulates low-dimensional embedding information sufficient to reconstruct complex geometric contours, such as high-curvature leading-edge details, with high fidelity.

Because the CNN encoder weights are frozen during the Cheby-KAN regression training, the framework does not undergo joint, end-to-end optimization. Consequently, the extracted latent space is strictly optimized for geometric reconstruction fidelity (SDF minimization) rather than being explicitly tailored for aerodynamic pressure regression (minimization of C_{p} prediction error). From an optimization standpoint, an end-to-end coupled training approach might discover a theoretically superior latent manifold that directly correlates geometric perturbations with localized pressure gradients.

However, this decoupled two-stage strategy was deliberately adopted to balance computational efficiency and geometric representation integrity. Jointly training the deep ResNet-based CNN alongside the Cheby-KAN on a massive dataset of 8.9×10^7 samples would incur prohibitive GPU memory overhead and drastically increase training times. Furthermore, isolating the CNN training prevents the phenomenon of representation collapse, ensuring that $\boldsymbol{z}_{\text{geo}}$ acts as a mathematically stable, purely geometric descriptor independent of varying flow conditions (Re, α). Therefore, while the two-stage pipeline sacrifices strict end-to-end theoretical optimality, it provides a highly stable, modular, and computationally tractable framework for large-scale aerodynamic surrogate modeling.

To model the strong nonlinearities and high-gradient flow patterns inherent in the massive dataset (8.9×10^7, 20-dimensional data points) while maintaining the high parameter efficiency of the KAN framework, a deep Cheby-KAN structure was constructed. The network comprises six hidden layers, with the structural layout (neuron count) specified as follows

20—16—16—16—16—16—16—1

Although a layer width of 16 appears compact compared to standard MLP, each "edge" connecting node i in layer l to node j in layer $l+1$ within the Cheby-KAN architecture contains a learnable univariate activation function $\phi_{j,i}(x)$. In this study, these univariate functions are expanded into a shifted Chebyshev polynomial series of order $N=5$.

Through the composite transformations of the six hidden layers, the network culminates in an output layer consisting of a single neuron that generates the final scalar prediction. The prediction target is defined as the normalized pressure coefficient field (C_{p} field), shown as

$$\text{Output}=C_{\text{p}}(x,y|Re,\alpha,\text{Geometry}) \tag{22}$$

This single-scalar output design aligns with the physical normalization strategy detailed in the data preparation phase. It ensures that the model captures the universal nonlinear mapping between geometric morphology and flow physics, providing a scientific foundation for its generalization across diverse, unseen airfoil configurations.

All experiments were conducted on the same hardware platform to ensure a fair comparison between the proposed Cheby-KAN and the baseline MLP. The training was performed on a workstation equipped with an NVIDIA RTX 4070 GPU (12 GB memory) and 32 GB RAM. Both models were implemented using PyTorch 2.5.1 under Python 3.12 with CUDA acceleration enabled.

4 Results

4.1 Baseline comparison and interpolation capability

In this study, a conventional MLP was selected as the primary baseline because it is the most widely adopted regression architecture for data-driven surrogate modeling in aerodynamic prediction. Mathematically, KAN is conceptualized as a direct alternative to traditional MLP by transitioning the learnable activation functions from nodes to edges. Therefore, a direct head-to-head comparison with an MLP of equivalent parameter scale serves as the most fundamental and controlled benchmark to isolate and verify the intrinsic representation capability of the Chebyshev polynomial basis functions without confounding structural factors.

Both Cheby-KAN and MLP were trained using the identical optimizer (Adam), the same initial learning rate (1×10^{-3}), identical batch size, and the same learning-rate decay schedule. The maximum number of training epochs was fixed at 500 for both models. Training was terminated when the validation MSE no longer improved or when the maximum number of epochs was reached.

Fig.12 illustrates the evolution of training and validation losses for the Cheby-KAN and the baseline MLP model (comprising 10 hidden layers with 100 nodes each). The experimental results demonstrate that Cheby-KAN exhibits significant advantages in terms of convergence rate, final predictive accuracy, and numerical stability.

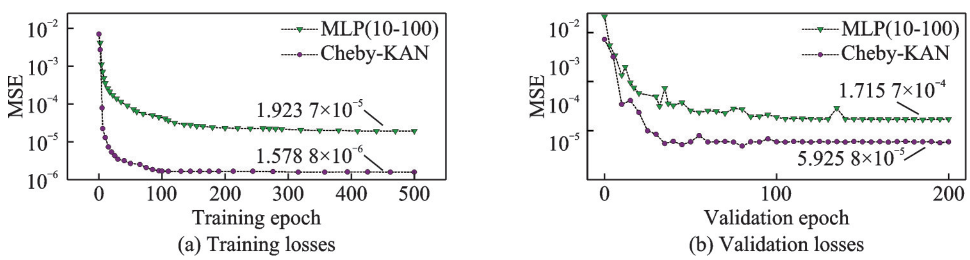

Fig.12 Evolution of training and validation losses for various models

Regarding convergence extrema, the training MSE of Cheby-KAN reaches 1.5788×10^{-6}, representing a 92% reduction compared to the MLP's 1.9237×10^{-5}. Concurrently, the validation MSE is reduced by 65%. This performance leap stems from the strong synergistic effect between the shifted Chebyshev polynomial basis functions and the KAN architecture. Unlike the global activation mechanism inherent in MLP, Cheby-KAN achieves precise capture of local high-gradient flow field features through learnable edge functions, thereby significantly accelerating the optimization process.

Table 3 further evaluates model complexity and computational efficiency. Training time was measured on the same hardware platform using identical optimization settings and convergence criteria for both models. When processing a large-scale dataset of 8.9×10^7 samples, Cheby-KAN demonstrates exceptional efficiency: Its training duration is less than 2 h, approximately 60% of the time required by the MLP (4 h and 58 min). Furthermore, Cheby-KAN achieves a 77% reduction in mean absolute error (MAE) on the validation set, with an improvement in the coefficient of determination (R^2) exceeding 10%. These results validate that Cheby-KAN possesses a superior capacity to handle large-scale heterogeneous aerodynamic data while maintaining high parameter efficiency.

Fig.13 and Fig.14 present the pressure distribution contours and corresponding absolute error distributions in the near-wall region of the e387 and NACA0012 airfoils under unseen flow conditions,

Table 3 Performance comparison between MLP and Cheby-KAN

Model architecture	Parameter	Test MSE	MAE	R^2	Training time
MLP (Baseline)	93 000	1.7157×10^{-4}	0.031 76	0.987 9	4 h 58 min
Cheby-KAN (Ours)	1 936	5.9258×10^{-5}	0.007 28	>0.999	1 h 52 min

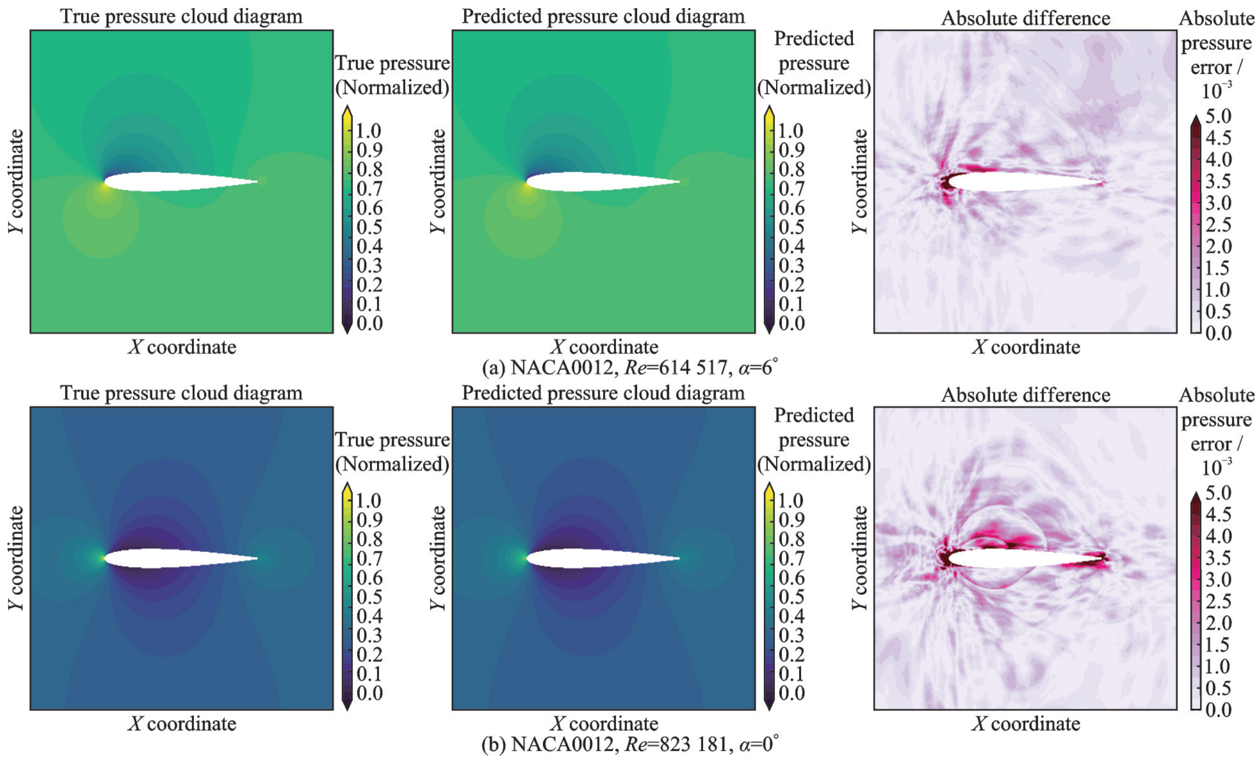

Fig.13 Flow field prediction results for the NACA0012 airfoil under various operational conditions

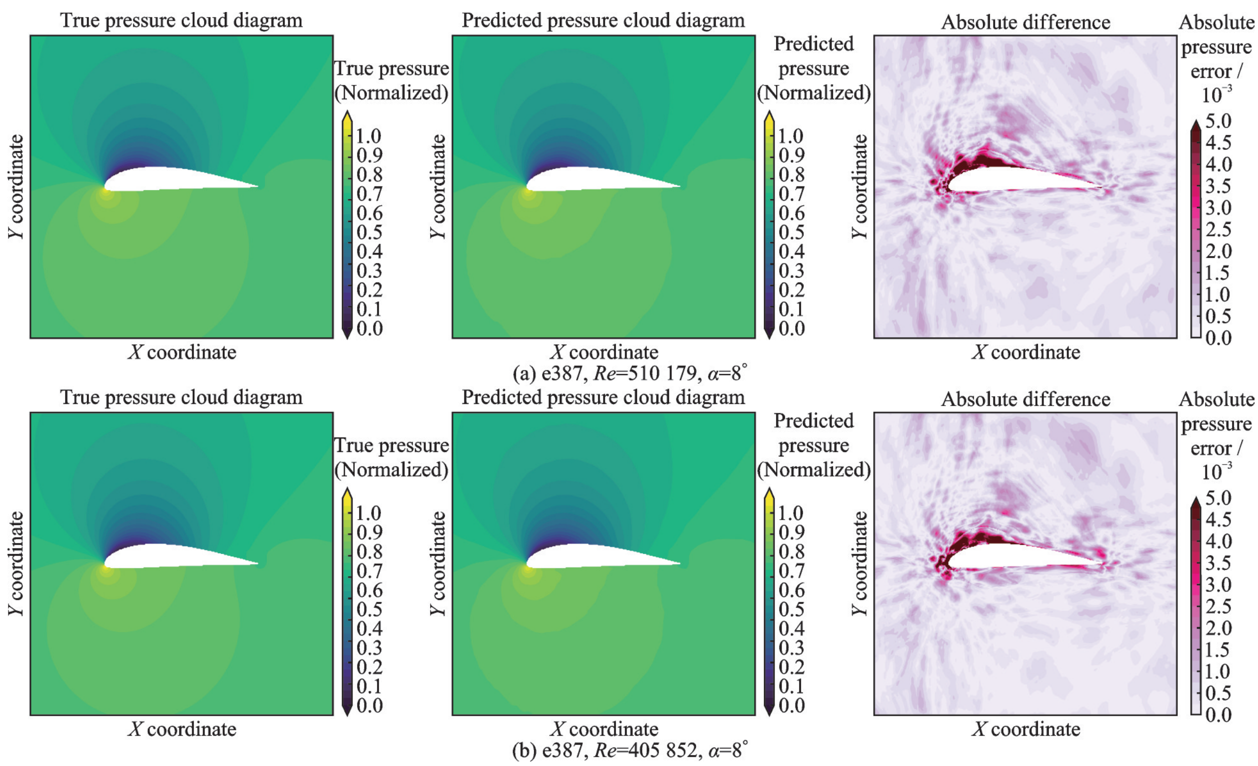

Fig.14 Flow field prediction results for the e387 airfoil under various operational conditions

aimed at validating the interpolation capability of the model. The results indicate that the maximum absolute error of the predicted pressure is strictly confined within 0.008, demonstrating exceptional physical fidelity.

The error peaks are primarily concentrated at the leading and trailing edges of the airfoil, and this distribution is closely associated with the underlying local flow physics rather than occurring randomly. Near the leading edge, the incoming flow decelerates to form a stagnation point before accelerating rapidly along the airfoil surface, producing extremely large pressure gradients. Consequently, even minor deviations in the predicted pressure field may be amplified into noticeable local errors. At the trailing edge, the interaction between the upper and lower boundary layers, together with the pressure recovery process and wake development, generates highly localized flow structures and sharp pressure variations, making this region inherently more difficult for surrogate models to approximate accurately.

Crucially, this physical correlation also reveals the architectural advantage of the Cheby-KAN framework over conventional MLP. The global activation mechanism employed by conventional MLP tends to suffer from error pollution, where the steep pressure gradients near the leading edge induce artificial oscillatory responses that propagate downstream, thereby deteriorating prediction accuracy in neighboring regions. In contrast, the Cheby-KAN framework effectively isolates these localized singular zones by exploiting the equioscillation property and rapid convergence of shifted Chebyshev polynomials. This localized representation enables the network to capture sharp pressure gradients while suppressing nonphysical oscillations, thereby preserving both the leading-edge suction peak and the trailing-edge pressure recovery slope without introducing excessive smoothing.

As illustrated by the comparative analysis of the e387 airfoil in Fig.15, Cheby-KAN substantially contracts the high-error envelopes at both the leading and trailing edges, achieving an approximately 50% reduction in maximum absolute error compared with the conventional MLP. These results demonstrate that the superior accuracy of Cheby-KAN originates not only from its enhanced approximation capability but also from its ability to accommodate the localized flow physics associated with stagnation, rapid acceleration, boundary-layer interaction, and wake evolution. Consequently, despite the intense pressure expansion near the leading edge, the residual errors remain well within the acceptable tolerance required for aerodynamic engineering design.

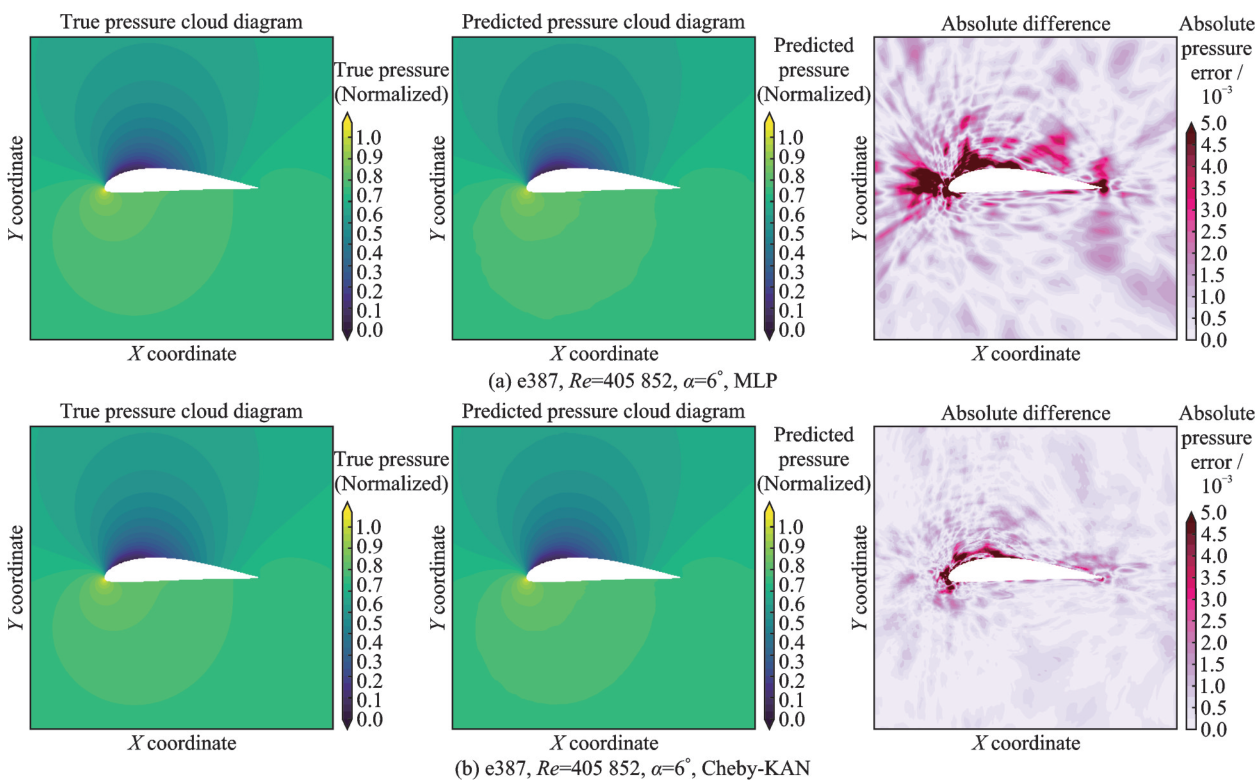

Fig.15 Flow field prediction results for the e387 airfoil between two different models

4.2 Validation of physical consistency

To evaluate the physical consistency of the surrogate model, this section performs a quantitative validation against key aerodynamic quantities derived from the CFD simulations.

Fig.16 illustrates the correlation between the lift coefficients predicted by Cheby-KAN and the ground-truth CFD data. The predictions for all test samples fall strictly within the 90%—110% error bands, demonstrating exceptional linear correlation and predictive stability.

Fig.17 provides a comparative analysis of the pressure coefficient curves across various angles of attack. Cheby-KAN not only accurately recovers the magnitude of the pressure coefficients but also precisely captures the location of suction peaks and the slope of the pressure recovery region. This fidelity in capturing C_p distribution details ensures that the model can provide reliable gradient guidance for subsequent aerodynamic optimization tasks, underscoring its scientific value as a high-fidelity surrogate model.

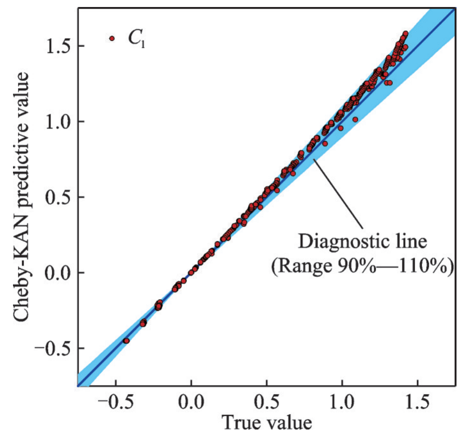

Fig.16 Comparison between predicted and ground-truth lift coefficients

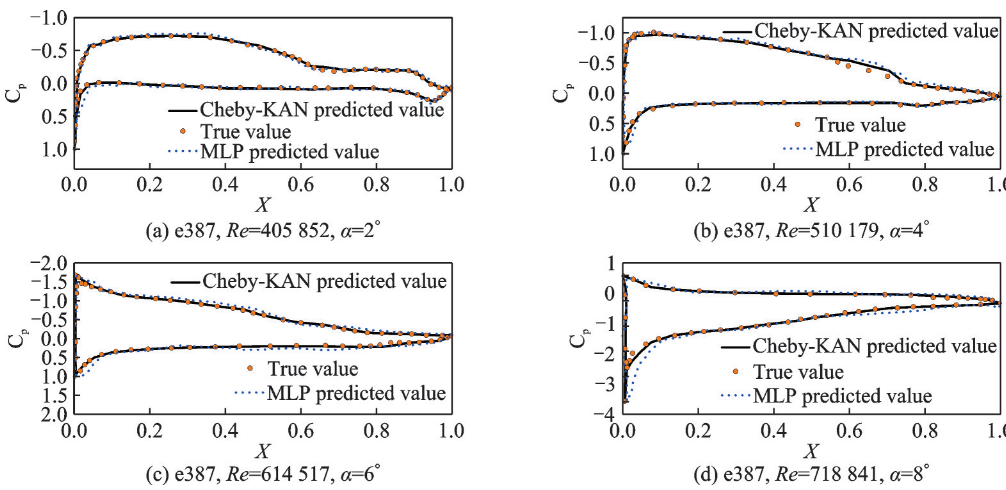

Fig.17 Prediction of C_p for the e387 airfoil under diverse operational conditions using the Cheby-KAN model

The high degree of agreement between the predicted and CFD-derived results indicates that the proposed model can accurately reproduce the integral aerodynamic quantities associated with the pressure distribution, such as the lift coefficient, within the investigated operating conditions. Furthermore, comparative analysis reveals that while conventional MLP often exhibits numerical oscillations in high-gradient regions near the leading edge, Cheby-KAN maintains extremely smooth yet sharp edge reconstruction characteristics. This not only proves the superior nonlinear fitting capability of the architecture but also reflects its effective suppression of the Runge phenomenon, ensuring predictive stability in the vicinity of geometric singularities.

4.3 Generalization performance validation

To rigorously evaluate the universality of the proposed framework, this study introduces the mh38 airfoil, which was entirely excluded during the training phase, for a "zero-shot" performance assessment. This test aims to examine whether the coupling between the CNN-extracted low-dimensional latent features and the Cheby-KAN regressor can accurately capture transferable aerodynamic relationships across unseen airfoil geometries, rather than merely memorizing specific geometric patterns encountered during training.

Fig.18 presents the pressure distribution contours and the corresponding absolute error field for the MH38 airfoil under the flow condition of $Re=$

6.0×10^5 and $\alpha=4°$. The results indicate that although the model had no prior exposure to the labeled data of this specific airfoil, the reconstructed pressure field maintains high consistency with the CFD ground truth. The MAE across the entire field is only 2.1×10^{-3}, with the maximum local error (max error) confined within 0.01, primarily localized in a small region near the leading edge.

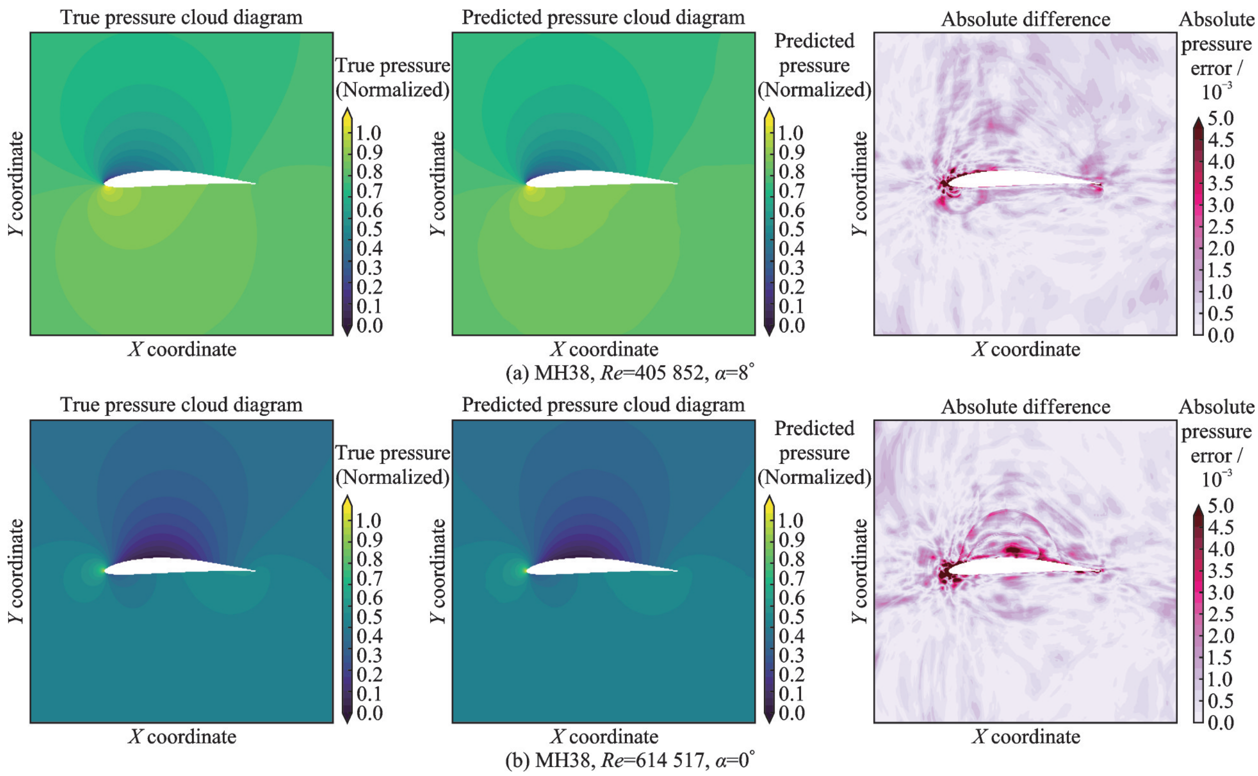

Fig.18 Flow field prediction results for the MH38 airfoil under various operational conditions

Fig.19 summarizes the distribution of maximum absolute errors across the entire operational envelope for three training airfoils (RAE2822, NACA0012, and e387) and three entirely unseen test airfoils (HD45, MH43, and MH64). It is observed that although the absolute errors for the test airfoils exhibit a marginal increase, no significant performance decay occurs. This provides compelling evidence that the integration of SDF-based geometric representation with the potent functional approximation capa-

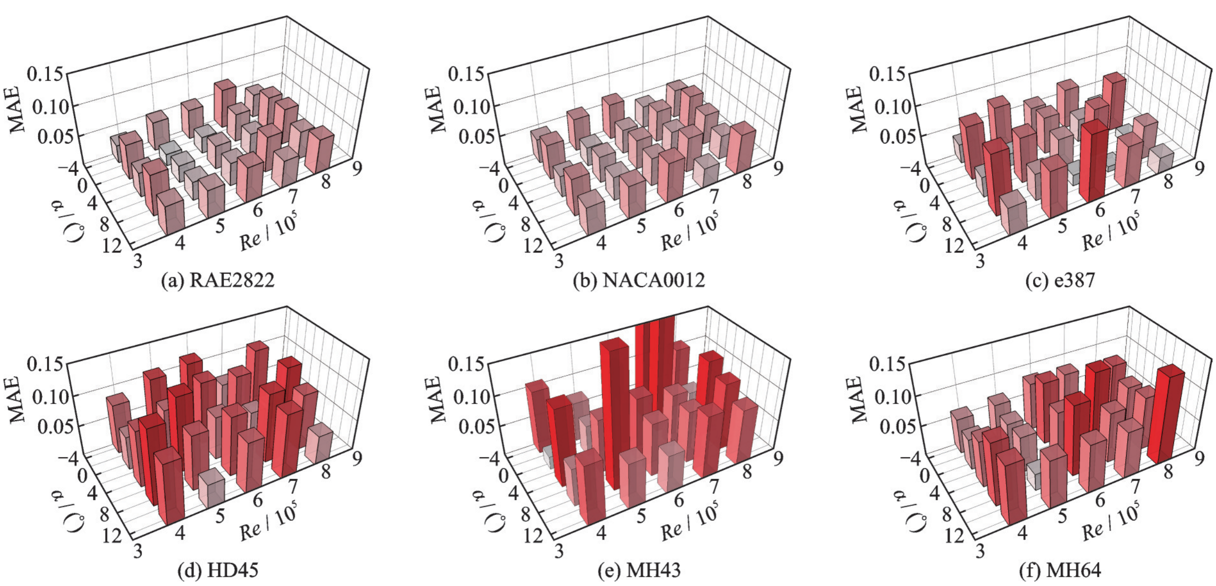

Fig.19 Distribution of maximum absolute errors in pressure field predictions using the Cheby-KAN model

bility of the KAN network enables the model to achieve cross-geometry feature transfer.

To further explore the boundaries of the model's physical robustness, extrapolation tests were conducted on the SD2030 and SD8020 airfoils under extreme conditions beyond the training envelope. The experimental setup involved a flow velocity of 85 m/s (13% beyond the training upper limit), with α extended from -6° to 16°.

Fig.20 illustrates the predicted lift coefficient curves under these conditions. The results demonstrate that, even in the high-AOA regime approaching stall, the proposed model accurately reproduces the nonlinear variation of the lift coefficient. The maximum absolute error between the predicted C_l and the CFD reference remains below 0.05. These results indicate that the proposed Cheby-KAN framework exhibits strong extrapolation capability beyond the training envelope while maintaining good agreement with the corresponding CFD-derived aerodynamic quantities. This suggests that the learned nonlinear mapping remains robust for previously unseen flow conditions within the investigated range.

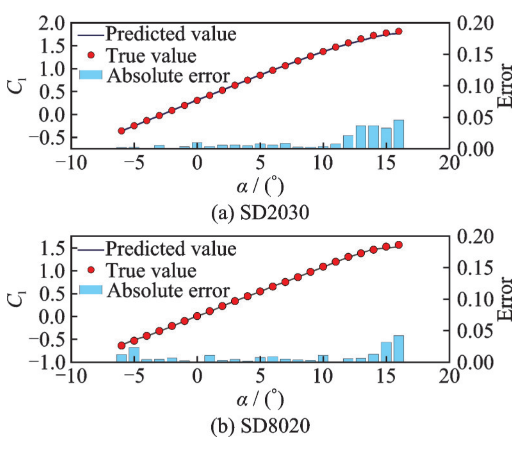

Fig.20 Comparison of predicted and CFD-calculated lift coefficients for SD2030 and SD8020 airfoils at 85 m/s

Although the extrapolation results indicate encouraging robustness beyond the training domain, these tests remain limited to conditions approaching stall rather than fully developed post-stall flows. Under massively separated flow conditions, the aerodynamic behavior is increasingly dominated by large-scale vortex shedding and strong nonlinear unsteady phenomena that are not represented in the present training database. Consequently, additional high-fidelity CFD or experimental datasets covering deep-stall conditions will be required before extending the proposed framework to such applications.

5 Conclusions

In this study, a novel deep learning framework, CNN-Cheby-KAN, is proposed for the high-fidelity and efficient prediction of two-dimensional airfoil pressure distributions. By integrating the automatic geometric feature extraction of CNN with the potent functional approximation of Cheby-KAN, the model addresses the critical bottlenecks of traditional surrogate models in terms of accuracy and computational cost. The main conclusions are summarized as follows.

(1) Superior predictive performance and efficiency: Experimental results demonstrate that the CNN-Cheby-KAN model significantly outperforms conventional MLP. Specifically, the training MSE is reduced by 92% (to 1.5788×10^{-6}), and the MAE on the validation set is decreased by 77%. Notably, the proposed model achieves these gains with a 98% reduction in parameters (1 936 vs. 93 000) and a 60% faster training convergence (less than 2 h) compared to the baseline MLP.

(2) Robustness in high-gradient regions: The introduction of shifted Chebyshev polynomials as basis functions effectively suppresses the Runge phenomenon, ensuring numerical stability near geometric singularities. The model maintains exceptional fidelity at the leading-edge stagnation point and trailing edge, with the maximum absolute pressure error strictly confined within 0.008 for unseen conditions.

(3) Strong generalization and physical consistency: The framework exhibits remarkable "zero-shot" generalization capability on completely unseen airfoil geometries (e.g., MH38), maintaining a field-wide MAE of only 2.1×10^{-3}. Furthermore, the high correlation between predicted lift coefficients and CFD ground truth (within 10% error bands) confirms that the predicted pressure distribu-

tions remain highly consistent with the corresponding integral aerodynamic quantities derived from CFD.

(4) Engineering utility: By bypassing the need for manual geometric parameterization (e.g., CST or NURBS) through SDF-based encoding, the CNN-Cheby-KAN framework provides a robust and automated tool for rapid aerodynamic analysis and optimization, offering a high-fidelity alternative to computationally expensive CFD simulations in early-stage aerospace design.

Although the proposed hybrid framework achieves the high precision, its training data are confined to steady, pre-stall regimes ($4\times10^5 \leqslant Re \leqslant 8\times10^5$, $-4.5^\circ \leqslant \alpha \leqslant 11.5^\circ$), causing performance to degrade under deep stall or massive flow separation. Additionally, the currently adopted decoupled two-stage training pipeline sacrifices strict end-to-end optimality, and the framework remains restricted to 2D configurations. To overcome these limitations, future research will focus on expanding the dataset to incorporate unsteady, deeply separated flows and exploring joint end-to-end optimization. Finally, the hybrid architecture will be extended to 3D aerodynamic applications to meet practical industrial engineering demands.

References

[1] CAO Huajun, LI Hongcheng, ZENG Dan, et al. The state-of-art and future development strategies of green manufacturing[J]. China Mechanical Engineering, 2020, 31(2): 134-144. (in Chinese)

[2] BHATNAGAR S, AFSHAR Y, PAN S, et al. Prediction of aerodynamic flow fields using convolutional neural networks[J]. Computational Mechanics, 2019, 64(2): 525-545.

[3] TAN Jianfeng. Influence of helicopter rotor on tail rotor unsteady aerodynamic loads[J]. Acta Aeronautica et Astronautica Sinica, 2015, 36(10): 3228-3240. (in Chinese)

[4] VEERS P, BOTTASSO C L, MANUEL L, et al. Grand challenges in the design, manufacture, and operation of future wind turbine systems[J]. Wind Energy Science, 2023, 8(7): 1071-1131.

[5] BLAZEK J. Computational fluid dynamics: Principles and applications[M]. 3rd ed. Oxford: Butterworth-Heinemann, 2015.

[6] TIWARI G, KUMAR J, PRASAD V, et al. Utility of CFD in the design and performance analysis of hydraulic turbines—A review[J]. Energy Reports, 2020, 6: 2410-2429.

[7] TOLLE K M, TANSLEY D S W, HEY A J. The fourth paradigm: Data-intensive scientific discovery[J]. Proceedings of IEEE, 2011, 99(8): 1334-1337.

[8] BIKMUKHAMETOV T, JÄSCHKE J. Combining machine learning and process engineering physics towards enhanced accuracy and explainability of data-driven models[J]. Computers & Chemical Engineering, 2020, 138: 106834.

[9] ROZOV V, BREITSAMTER C. Data-driven prediction of unsteady pressure distributions based on deep learning[J]. Journal of Fluids and Structures, 2021, 104: 103316.

[10] OBIOLS-SALES O, VISHNU A, MALAYA N, et al. CFDNet: A deep learning-based accelerator for fluid simulations[C]//Proceedings of the 34th ACM International Conference on Supercomputing. Barcelona, SPAIN: Association for Computing Machinery, 2020: 1-12.

[11] MI B, CHENG W. Intelligent aerodynamic modelling method for steady/unsteady flow fields of airfoils driven by flow field images based on modified U-Net neural network[J]. Engineering Applications of Computational Fluid Mechanics, 2025, 19(1): 2440075.

[12] THUEREY N, WEIßENOW K, PRANTL L, et al. Deep learning methods for Reynolds-averaged Navier-Stokes simulations of airfoil flows[J]. AIAA Journal, 2020, 58(1): 25-36.

[13] TOMPSON J, SCHLACHTER K, SPRECHMANN P, et al. Accelerating Eulerian fluid simulation with convolutional networks[C]//Proceedings of the 34th International Conference on Machine Learning. Sydney, Australia: PMLR, 2017: 3421-3430.

[14] WANG Y, SHIMADA K, BARATI FARIMANI A. Airfoil GAN: Encoding and synthesizing airfoils for aerodynamic shape optimization[J]. Journal of Computational Design and Engineering, 2023, 10(4): 1350-1362.

[15] ESFAHANIAN V, IZADI M J, BASHI H, et al. Aerodynamic shape optimization of gas turbines: A deep learning surrogate model approach[J]. Structural and Multidisciplinary Optimization, 2024, 67(1): 2.

[16] ZHANG Zeyu, ZHOU Xuxi, XU Nan, et al. Advances in machine learning for wind-induced fluid-structure interaction of large-scale structures[J]. Acta Aerody-

namica Sinica, 2025, 43(5): 53-77. (in Chinese)

[17] SEKAR V, JIANG Q, SHU C, et al. Fast flow field prediction over airfoils using deep learning approach[J]. Physics of Fluids, 2019, 31(5): 057103.

[18] LIU Z, WANG Y, VAIDYA S, et al. Kan: Kolmogorov-arnold networks[C]//Proceedings of the International Conference on Learning Representations. Singapore: OpenReview, 2025: 70367-70413.

[19] LECUN Y, BOTTOU L, BENGIO Y, et al Gradient-based learning applied to document recognition[J]. Proceedings of IEEE, 2002, 86(11): 2278-2324.

[20] KRIZHEVSKY A, SUTSKEVER I, HINTON G E. Imagenet classification with deep convolutional neural networks[C]//Proceedings of Advances in Neural Information Processing Systems 25. Lake Tahoe, NV, USA: Curran Associates, Inc., 2012: 1097-1105.

[21] ZAFAR A, AAMIR M, MOHD NAWI N, et al. A comparison of pooling methods for convolutional neural networks[J]. Applied Sciences, 2022, 12(17): 8643.

[22] WENG L, ZHANG H, CHEN H, et al. Towards fast computation of certified robustness for ReLU networks[C]//Proceedings of the 35th International Conference on Machine Learning. Stockholmsmässan, Stockholm Sweden: PMLR, 2018: 5276-5285.

[23] SCHMIDT-HIEBER J. The Kolmogorov-Arnold representation theorem revisited[J]. Neural Networks, 2021, 137: 119-126.

[24] YEO S, NGUYEN P A, LE A N, et al. KAN-PDEs: A novel approach to solving partial differential equations using Kolmogorov-Arnold networks—Enhanced accuracy and efficiency[C]//Proceedings of International Conference on Electrical and Electronics Engineering. Singapore: Springer Nature, 2024: 43-62.

[25] KOENIG B C, KIM S, DENG S. KAN-ODEs: Kolmogorov-Arnold network ordinary differential equations for learning dynamical systems and hidden physics[J]. Computer Methods in Applied Mechanics and Engineering, 2024, 432: 117397.

[26] ZHANG Z, SHEN T, ZHANG Y, et al. AL-PKAN: A hybrid GRU-KAN network with augmented lagrangian function for solving PDEs[J]. Scientific Reports, 2025. DOI: 10.1038/S41598-025-92900-1.

[27] MAHMOUD A A, PESTER A, MUTTARDI M, et al. Cheby-KANs: Advanced Kolmogorov-Arnold networks for applying geometric deep learning in quantum chemistry applications[J]. IEEE Access, 2025, 13: 130525-130534.

[28] GUO C, SUN L, LI S, et al. Physics-informed Kolmogorov-Arnold network with Chebyshev polynomials for fluid mechanics[J]. Physics of Fluids, 2025. DOI: 10.48550/arXiv.2411.04516.

[29] AL-MEKHLAFI Z G, ALTMEMI J M, AL-SHAREEDA M A, et al. ChebIoD: A Chebyshev polynomial-based lightweight authentication scheme for internet of drones environments[J]. Scientific Reports, 2025, 15(1): 32897.

[30] SELIG M S. UIUC airfoil data site[EB/OL]. [2026-07-26]. https://m-selig.ae.illinois.edu/ads.html.

[31] YANG W, JIN W. ContactSDF: Signed distance functions as multi-contact models for dexterous manipulation[J]. IEEE Robotics and Automation Letters, 2025, 10(5): 4212-4219.

[32] NIU Z, LIU Q. Optical-SDFNet: A hybrid optical-driven neural operator for aerodynamic pressure field prediction[C]//Proceedings of the Second International Conference on Communication, Information, and Digital Technologies (CIDT 2025). Singapore: International Society for Optics and Photonics, 2026: 1066-1073.

[33] SELIG M S, MCGRANAHAN B D. Wind tunnel aerodynamic tests of six airfoils for use on small wind turbines[J]. Journal of Solar Energy Engineering, 2004, 126(4): 986-1001.

[34] SHIN J, BOND T. Results of an icing test on a NACA 0012 airfoil in the NASA lewis icing research tunnel[C]//Proceedings of the 30th Aerospace Sciences Meeting and Exhibit. Reno, NV, USA: American Institute of Aeronautics and Astronautics, 1992: 647.

[35] MCGHEE R J. Experimental results for the Eppler 387 airfoil at low Reynolds numbers in the Langley low-turbulence pressure tunnel[M]. Washington D C, USA: National Aeronautics and Space Administration, 1988.

[36] CRITZOS C C, HEYSON H H, BOSWINKLE JR R W. Aerodynamic characteristics of NACA 0012 airfoil section at angles of attack from 0 deg to 180 deg: NACA-TN-3361[R]. [S.l.]: National Aeronautics and Space Administration, 1955: 1-24.

Acknowledgement This work was supported by the the Yangzhou Science and Technology Plan Project (No. YZ2025232).

Authors

The first author Mr. **CHEN Yaohong** received his Bachelor's degree in engineering from Northwestern Polytechnical University in 2023, and is currently pursuing a Master's degree at AVIC Aerodynamics Research Institute. His research interests include deep learning, aerodynamic modeling, and computational fluid dynamics (CFD) simulation.

The corresponding author Prof. **JIAO Renshan** is a senior engineer at the AVIC Aerodynamics Research Institute and a first level technical expert of AVIC Group. His research interests include wind tunnel testing techniques, wall interference correction, and support interference correction.

Author contributions Mr. **CHEN Yaohong** conceived the study, developed the methodology, implemented the deep-learning models, and analyzed the results. Mr. **ZHANG Luchi** contributed to methodology development, model training, and validation. Prof. **DENG Yiju** and Prof. **YU Yanze** contributed to technical review. Ms. **LI Xiang** contributed to CFD simulation design and result interpretation. Prof. **JIAO Renshan** supervised the research, reviewed the manuscript, and provided overall guidance. All authors reviewed and approved the final manuscript.

Competing interest The authors declare no competing interests.

(Production Editor: ZHANG Huangqun)

基于混合CNN与Chebyshev-Kolmogorov-Arnold网络的高精度翼型压力分布预测模型

陈耀鸿[1],张路驰[2],邓一菊[3],于彦泽[3],李　响[4],焦仁山[1]

(1.中国航空工业空气动力研究院哈尔滨分院,哈尔滨 150001,中国;2.中国燃气涡轮研究院,成都 610599,中国;3.中国航空研究院研究生院,扬州 225111,中国;4.中国航空工业空气动力研究院沈阳分院,沈阳 110034,中国)

摘要:在翼型气动设计与优化中,传统计算流体力学(Computational fluid dynamics, CFD)求解Navier-Stokes方程虽能获得高精度压力分布,但网格离散求解流程计算开销大,开展大批量参数扫描时耗时极其突出。现有深度学习气动代理模型普遍存在前缘/后缘高梯度流场捕捉能力不足、训练收敛慢、容易出现数值振荡、跨翼型泛化性能有限等缺陷。针对上述问题,提出一种卷积神经网络(Convolutional neural network, CNN)与切比雪夫-柯尔莫哥洛夫-阿诺尔德网络(Chebyshev-Kolmogorov-Arnold network, Cheby-KAN)相融合的混合深度学习框架,实现二维翼型表面及流场压力分布的高精度快速预测。该框架采用残差CNN作为几何编码器,直接对翼型符号距离场(Signed distance field, SDF)网格进行特征学习,自动将复杂翼型几何转化为16维低维几何潜向量,不再依赖CST、NURBS等传统人工几何参数化手段。回归部分对KAN网络进行改进,采用平移切比雪夫多项式替代原始B样条作为边函数的基函数;借助平移切比雪夫多项式的正交性与等振荡逼近特性,有效抑制高次插值带来的Runge振荡现象,提升翼型前缘驻点、尾迹区等高梯度压力区域的数值稳定性。基于ANSYS Fluent结合SST k-ω湍流模型构建包含8 900万矢量样本点的大规模CFD数据集,并严格按算例层级划分数据集以杜绝数据泄露。对比实验结果表明:相较于同等任务下的传统多层感知机(Multi-layer perceptron, MLP),Cheby-KAN模块参数量下降98%,训练耗时减少60%,训练均方误差(Mean square error, MSE)降低92%。在完全未参与训练的翼型样本上开展零样本泛化测试,全场压力预测平均绝对误差(Mean absolute error, MAE)仅为2.1×10^{-3};由预测压力积分得到的升力系数与CFD基准结果误差控制在10%以内,具备良好的物理一致性。该模型能够快速输出流场压力系数场,可为翼型多目标气动优化提供高保真、低成本的数据驱动代理工具。

关键词:翼型压力分布;深度学习;卷积神经网络;切比雪夫多项式;柯尔莫哥洛夫-阿诺德网络;气动预测

研究亮点:

1.构建CNN-Cheby-KAN混合代理模型,面向二维翼型压力分布高精度快速预测,CNN编码器基于符号距离场SDF自动提取几何潜特征,摆脱CST、NURBS等人工几何参数化依赖。

2.将平移切比雪夫多项式替代B样条嵌入KAN架构得到Cheby-KAN,利用多项式正交性有效抑制Runge振荡,显著提升翼型前缘、后缘等高梯度压力区域的数值稳定性。

3.实现未见翼型零样本泛化预测,未知翼型全场MAE仅为2.1×10^{-3};预测升力系数与CFD基准偏差控制在10%以内,具备工程气动优化的应用潜力。