

Integrated Aerodynamic and Stealth Design of Wing Airfoil Based on Generative Model

CHI Xinyan¹, ZENG Lifang^{1,2*}, LI Jun^{1,3}, LI Yuhang¹

1. School of Aeronautics and Astronautics, Zhejiang University, Hangzhou 310027, P. R. China; 2. Zhejiang Key Laboratory of Intelligent High Speed UAV Technology, Zhejiang University, Hangzhou 310027, P. R. China;
3. Huanjiang Laboratory, Zhuji 311800, P. R. China

(Received 28 November 2025; revised 16 March 2026; accepted 24 June 2026)

Abstract: Generative design methods have been widely applied in modern aircraft aerodynamic design. However, the integrated optimization of aerodynamic and stealth performance in aircraft still relies on surrogate models and multi-objective optimization algorithms. To address the complex verification and optimization procedures in current integrated aerodynamic-stealth aircraft design, this paper proposes a rapid generative design method based on a conditional denoising diffusion probability model (CDDPM). First, the class-shape transformation (CST) method is employed for parametric modeling of airfoils. To build the aerodynamic and stealth performance datasets, the vortex lattice method and the physical optics method for large-sized objects are used to compute the lift-to-drag ratio (L/D) and radar cross-section (RCS), respectively. Based on the dataset, a generative conditional diffusion model is implemented to achieve the mapping relationship from target performance (L/D and RCS) to CST parameters of wing airfoils. Validation results indicate that the prediction errors for the generative model in aerodynamic-stealth performance are smaller than 6%. Meanwhile, the generated airfoils exhibit notable diversity. Furthermore, optimization design of airfoils considering both aerodynamic and stealth performance is conducted, where the diffusion model is utilized to generate new airfoils to expand the design space. The pareto front is obviously expanded with the minimum RCS decreased by 28.6%, and the maximum L/D increased by 7.5%. This study establishes a generative model-based framework for rapid aerodynamic-stealth optimization of airfoils, laying a foundation for AI-driven multidisciplinary design optimization (MDO) in aircraft design.

Key words: integrated aerodynamic-stealth design; generative design; diffusion model; class-shape transformation (CST) parameterization; lift-to-drag ratio; radar cross-section (RCS)

CLC number: V211.41

Document code: A

Article ID: 1005-1120(2026)03-0427-16

0 Introduction

The rapid advancement of modern air defense systems, particularly in radar detection and missile interception technologies, has posed severe threats to aircraft survivability. The traditional sequential design philosophy of “aerodynamics first, stealth second” can no longer satisfy the stringent requirements for aircraft survivability and combat effectiveness in next-generation high-tech localized conflicts. In this context, the integrated design of aerodynamic and stealth performance has emerged as a core

methodology for enhancing the overall performance of aircraft, representing an inevitable trend and a cutting-edge direction in the development of advanced aerospace equipment. Over the past decade, researchers have frequently adopted aerodynamic-stealth integrated design methods based on collaborative optimization using surrogate models^[1] and multi-objective optimization algorithms^[2]. However, such approaches, which rely on iterative simulations or gradient-based optimization^[3], are often hindered by high computational costs and a tendency to converge to local optima. Although surrogate mod-

*Corresponding author, E-mail address: lifang_zeng@zju.edu.cn.

How to cite this article: CHI Xinyan, ZENG Lifang, LI Jun, et al. Integrated aerodynamic and stealth design of wing airfoil based on generative model[J]. Transactions of Nanjing University of Aeronautics and Astronautics, 2026, 43(3): 427-442.

<http://dx.doi.org/10.16356/j.1005-1120.2026.03.008>

els can expedite the evaluation of physical properties, they encounter the “curse of dimensionality”^[4-5] in high-dimensional design spaces, which severely restricts their applicability in engineering practice.

In recent years, the breakthrough progress in deep learning^[6] across fields such as natural language processing^[7] and computer vision^[8], along with its successful application in scientific computing^[9], has opened up new opportunities for complex engineering optimization. Among various deep learning approaches, conditional diffusional models (CDMs)^[10]—capable of generating samples that satisfy specific constraints—exhibit a mechanism that closely mirrors the inverse design process in aircraft design: Directly producing geometric configurations based on target performance requirements. The successful deployment of large-scale conditional generative models, such as those for text-to-image^[11] and text-to-video^[12] generation, fully demonstrates the considerable potential of such technologies in aircraft design under complex multidisciplinary constraints. Within the family of generative models, denoising diffusion probabilistic models (DDPMs)^[13] have attracted growing attention due to their training stability, effective mitigation of issues such as mode collapse^[14] and posterior collapse^[15], as well as their demonstrated higher optimization efficiency and lower feasibility error in topology optimization tasks^[16].

In the field of airfoil design, research incorporating diffusion models has primarily evolved along three directions. First, multidisciplinary conditional control generation: for instance, Gong et al.^[17] proposed a latent diffusion model framework integrated with a transformer autoencoder to extract latent representations of airfoil point clouds, enabling conditional diffusion generation in the latent space under multi-objective constraints such as aerodynamic and geometric characteristics. However, the generated airfoils suffered from a lack of smoothness. Wang et al.^[18] employed CDM to generate airfoils conditioned on buffet lift coefficient, cruise drag coefficient, and thickness, introducing classifier-free guidance to enhance conditional control capability. Cheng et al.^[19] investigated a fast inverse design

method for airfoils based on conditional denoising diffusion probability model (CDDPM), using airfoil coordinates directly as the shape representation to achieve rapid design under aerodynamic performance constraints. Second, parametric latent space construction: Zhang et al.^[20] combined the class-shape transformation (CST) parameterization, an autoencoder (AE), and a denoising diffusion implicit model (DDIM) to construct a latent diffusion model (LDM). This approach generates airfoils in the latent space, where AE reduces the dimensionality of CST parameters to latent variables, and DDIM learns their distribution—effectively mitigating the “curse of dimensionality”. Tao et al.^[21] utilized a DDPM to generate pressure distributions, which were then mapped to CST parameters. Third, real-time optimization attempts in geometric space: Graves et al.^[22] built a CDM directly in geometric space (without a latent representation) to generate airfoils conditioned on lift or drag coefficients. Principal component analysis (PCA) revealed that the model explored regions of the design space not covered by the training set^[22]. Yonekura et al.^[23] explored the application of DDPM in airfoil generation and observed that Gaussian noise introduced during the diffusion process led to non-smooth shapes, and they proposed moving average filtering as a smoothing solution.

Although the aforementioned studies have made significant progress in airfoil generation, several key limitations remain. First, most existing works primarily consider aerodynamic coefficients or airfoil surface pressure distributions as constraints for generative design, overlooking the critical aspect of stealth performance. In practice, aerodynamic and stealth objectives often conflict, necessitating dynamic trade-off strategies during design. Second, the engineering adaptability of current approaches is limited. Representing airfoils via discrete point clouds lack the global control capability offered by parametric methods, making it difficult to satisfy airfoil smoothness requirements^[24]. Third, while generative model-based airfoil generation has advanced, it typically does not fully replace the traditional airfoil optimization process. Conversely, tradi-

tional inverse design methods often rely on discriminant models^[25], which restricts the search to the original training set and impedes effective optimization. Therefore, a major challenge lies in effectively integrating these two design paradigms to overcome their respective shortcomings.

To bridge this gap, this paper proposes an aerodynamic-stealth integrated optimization design framework for airfoils based on CDDPM. The main contributions of this work are summarized as follows.

(1) A new problem paradigm: We pioneer the application of generative diffusion models to the coupled aerodynamic-stealth design problem, demonstrating that a single data-driven model can capture the complex trade-off between lift-to-drag ratio (L/D) and radar cross-section (RCS).

(2) A novel conditional generative framework: We introduce a joint conditional injection mechanism that incorporates both L/D and RCS targets into the diffusion process, enabling direct and accurate mapping from performance specifications to airfoil CST parameters. The framework establishes a complete process encompassing automated modeling, simulation, and generative design.

(3) Superior design space exploration: Unlike

traditional surrogate-based optimization that is confined to the training set, our CDDPM not only reproduces existing designs with high fidelity (errors $< 6\%$) but also actively generates novel airfoils that lie beyond the original Pareto front.

This paper is organized as follows. Section 1 presents the methodological framework, including the CST parameterization, aerodynamic and stealth simulation methods, and the training strategy of CDDPM. Section 2 provides a comprehensive analysis of the generative design and discusses the optimization results. Finally, Section 3 draws a conclusion for this work.

1 Diffusion Model-Driven Method for Airfoil Generative Design

The aerodynamic-stealth integrated optimization design method for airfoils based on a generative model, as developed in this study, comprises four key steps: Parametric representation of airfoils, construction of an aerodynamic-stealth performance dataset, performance-to-shape mapping via CDPM, and design space expansion. The overall workflow is illustrated in Fig.1, and the specific technical procedure is outlined as follows.




Fig.1 Framework of the generative optimization design method for airfoils

(1) **Parametric airfoil representation:** The CST method^[26] is employed for efficient geometric representation of airfoils. Combined with latin hypercube sampling (LHS)^[27], a variety of airfoil shape samples are generated to establish a baseline dataset encompassing both aerodynamic and stealth performance attributes.

(2) **Multidisciplinary performance dataset construction:** Using the VSPAero Simulation platform based on vortex lattice method (VLM)^[28] for aerodynamic analysis and the Altair Feko Simulation platform based on physical optics (PO) method for stealth analysis^[29], the lift-to-drag ratio (L/D) and RCS of the generated airfoil samples are systematically computed. This process yields a high-quality, multidisciplinary performance dataset.

(3) **Conditional generative modeling:** CD-DPM is constructed to learn the probability distribution of CST shape parameters under specified target performance conditions (L/D and RCS). Through a stochastic denoising process, the model generates novel airfoil shapes that satisfy the performance requirements while maintaining shape diversity.

(4) **Design space expansion and optimization:** The trained diffusion model is utilized to batch-generate new airfoil samples. Their performance is subsequently verified via simulation, thereby expanding the original design space and yielding optimized solutions that approach or exceed the original Pareto front. Ultimately, this establishes a comprehensive intelligent design framework for aerodynamic-stealth integration, covering parametric modeling, performance analysis, probabilistic generation, and design space augmentation.

In this study, the training dataset is constructed using the NACA 63006 airfoil as the baseline geometry. Based on this profile, a straight three-dimensional wing with a span of 2 m and a chord length of 1 m is modeled as the fixed configuration. In this setup, all geometric parameters—such as span and sweep angle—are held constant except for the sectional airfoil shape, which serves as the primary design variable. The CST parameterization

method is applied to represent and control the two-dimensional cross-sectional geometry. A total of 1 360 distinct airfoil profiles are generated via LHS. Aerodynamic and stealth performance simulations are performed on the corresponding 3D wing model to compute the L/D and RCS values for each sample. Finally, a diverse training dataset for the CD-DPM is constructed, where the airfoil shape parameters serve as generation targets and the corresponding performance metrics (L/D and RCS) act as conditioning inputs.

1.1 Dataset

1.1.1 Geometric shape data

This study employs the CST method with 20 control points (10 for the upper surface and 10 for the lower surface, corresponding to a polynomial order of $n = 9$) for parametric representation of airfoil geometry. To generate a diverse training dataset, LHS is applied to perturb the 19 independent CST parameters (the first lower surface control point is fixed as the negative of the first upper surface control point to ensure a sharp trailing edge) within a range of $\pm 50\%$ relative to the baseline NACA 63006 airfoil values. This perturbation range was chosen to ensure sufficient geometric diversity while maintaining physically realistic airfoil shapes. Through this process, 1 360 distinct airfoil profiles are generated. Fig.2 visualizes the distribution of the LHS-perturbed samples, illustrating the geometric coverage achieved.




Fig.2 LHS perturbed airfoil samples

1.1.2 Aerodynamic-stealth performance data

Aerodynamic performance is evaluated under the following flow conditions: Mach number of 0.8, Reynolds number of 10 million, and an angle of attack of 2° . Numerical simulations are carried out using VLM, which is based on potential flow theory. The governing equation is the Laplace equation, shown as

$$\nabla^2 \phi = 0 \quad (1)$$

where ϕ is the velocity potential function.

The wall boundary condition satisfies

$$\left. \frac{\partial \phi}{\partial n} \right|_{\text{wall}} = -U_{\infty} \cdot \hat{n} \quad (2)$$

where $\hat{n}$ is the unit normal vector to the wall surface and U_{∞} the freestream velocity.

The velocity potential is decomposed as

$$\phi = U_{\infty} \cdot \mathbf{r} + \sum_{j=1}^N \gamma_j \varphi_j(\mathbf{r}) \quad (3)$$

where γ_j is the vortex strength of the j th panel, φ_j the velocity potential function induced by the vortex core, N the total number of discretized panels on the surface, and $\mathbf{r}$ the position vector of the field point.

An adaptive Cartesian grid is used, with refinement near the wall. Key aerodynamic parameters such as lift coefficient, drag coefficient, and lift-to-drag ratio L/D are calculated through this method.

To validate the accuracy of the computational method, the simulation results for the NACA0015 airfoil under the specified flow conditions were compared with experimental data. In Figs.3, 4, the lift coefficient C_l and drag coefficient C_d are the primary dimensionless aerodynamic force coefficients, and the angle of attack α denotes the angle between the chord line of the airfoil and the incoming flow direction. As shown in Fig.3 and Fig.4, under the operational conditions considered in this study, the calculated lift coefficient exhibits an error of 2.60%, the calculated drag coefficient exhibits an error of 2.86%, demonstrating that the adopted aerodynamic computation method provides sufficient accuracy for engineering applications^[30].




Fig.3 Comparison of calculated lift coefficient for NACA0015 airfoil with experimental data




Fig.4 Comparison of calculated drag coefficient for NACA0015 airfoil with experimental data

Stealth performance evaluation focuses on the RCS of the airfoil under X-band electromagnetic wave irradiation. The simulation conditions are set as: frequency $f = 10$ GHz, incident wave pitch angle range from -30° to 30° , horizontal polarization (HH). Numerical calculation is based on the PO method for large-sized objects. The induced current J_s is approximately expressed as different distributions in the illuminated and shadowed regions, shown as

$$J_s \approx \begin{cases} 2\hat{n} \times H^{\text{inc}} & \text{Irradiation region} \\ 0 & \text{Shadowed region} \end{cases} \quad (4)$$

where H^{inc} is the intensity of the incident magnetic field.

The scattered field integral expression is

$$E^{\text{scat}}(\mathbf{r}) = \frac{jk\eta_0}{4\pi} \int [J_s - (J_s \cdot \hat{r})\hat{r}] \frac{e^{-jkR}}{R} dS \quad (5)$$

where $k = 2\pi f/c$ is the wavenumber with c being the speed of light in vacuum, $\eta_0 = 377 \Omega$ the free space wave impedance, $R = |\mathbf{r} - \mathbf{r}'|$ the distance from the field point to the source point, and S the surface of scatterer.

The monostatic RCS is defined as

$$\sigma = \lim_{R \rightarrow \infty} 4\pi R^2 \frac{|E^{\text{scat}}|^2}{|E^{\text{inc}}|^2} \quad (6)$$

where σ is the radar cross-sectional area, i. e., RCS, in units of m^2 ; E^{scat} the intensity of the scattered electric field; and E^{inc} the intensity of the incident electric field.

To verify the correctness of the RCS calculation method, a flat plate structure was selected as the verification case, with a plate length and width of 1 m and a thickness of 5 mm. The theoretical ap-

proximate RCS value for the plate^[31] is

$$\sigma_p = \frac{4\pi A^2}{\lambda^2} \quad (7)$$

where A is the plate area, and λ the wavelength of electromagnetic wave, for electromagnetic wave normal incidence, i.e., incident angle 0° .

According to this formula, the theoretical RCS value of the plate at 0° incidence is 139.62 m^2 . Under normal incidence conditions, the relative error between the calculated result and the theoretical value is 1.55% , verifying the correctness of the stealth performance calculation method.

The final obtained dataset contains the CST parameters of 1 360 airfoil samples and their corresponding aerodynamic and stealth performance indicators. Fig.5 shows the distribution of dataset samples in the L/D -RCS condition space, reflecting good diversity and coverage of the samples, which can meet the training requirements of the subsequent conditional diffusion model.




Fig.5 Distribution of dataset samples in the L/D -RCS performance space

1.2 Denoising diffusion probabilistic model

DDPM is an important method that has achieved remarkable breakthroughs in the field of deep generative models in recent years. Its core idea originates from non-equilibrium thermodynamics, learning the data distribution by simulating a progressive data “destruction” (noising) and “reconstruction” (denoising) process.

1.2.1 Standard diffusion model mechanism

DDPM consists of two core processes: The forward diffusion process and the reverse diffusion process.

Forward diffusion process This process transforms the original data sample x_0 (e.g., a CST sequence of an airfoil) into noise x_T that approximately follows a standard normal distribution $N(0, I)$ through T discrete time steps. This is a Markov chain process, and its conditional probability distribution is defined as

$$q(x_t|x_{t-1}) = N(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t I) \quad (8)$$

where $t = 1, 2, \dots, T$; β_t is the pre-defined noise variance schedule, controlling the amount of noise added at each step (typically β_t increases with t).

Using the reparameterization trick, the noised state at any step t can be directly calculated from x_0 , i.e.

$$q(x_t|x_0) = N(x_t; \sqrt{\bar{\alpha}_t} x_0, (1 - \bar{\alpha}_t) I) \quad (9)$$

where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$. As t increases, $\bar{\alpha}_t \rightarrow 0$, x_t gradually loses the original data information, approaching pure noise.

The core of DDPM lies in learning the reverse process. The model trains a neural network $\epsilon_\theta(x_t, t)$, whose goal is to predict the noise ϵ added to x_{t-1} during the forward process, based on the current noised state x_t and time step t . The training objective is to minimize the mean squared error (MSE) between the predicted noise and the true noise, i.e.

$$\mathcal{L}_{\text{simple}} = E_{x_0, \epsilon \sim N(0, I), t} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2] \quad (10)$$

where $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$.

By optimizing this loss function, the network learns the reverse process of gradually removing noise and restoring the data structure.

Reverse diffusion process When generating new samples, an initial noise is sampled from the standard normal distribution $x_T \sim N(0, I)$. Then, using the trained noise prediction network ϵ_θ , denoising is performed step by step from $t = T$ to $t = 1$. At each step, the noise $\epsilon_\theta(x_t, t)$ is predicted based on the current state x_t , and the mean and variance for x_{t-1} are calculated as

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) \quad (11)$$

$$\sigma_t^2 = \frac{(1 - \alpha_t)(1 - \bar{\alpha}_{t-1})}{1 - \bar{\alpha}_t} \quad (12)$$

Finally x_{t-1} is sampled from the distribution $q(x_{t-1}|x_t) = N(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2 I)$. After T iterations, a clear data sample x_0 is obtained.

1.2.2 Conditional denoising diffusion probabilistic model

In the airfoil generative design task, our goal is to directly generate airfoil geometries that satisfy specific aerodynamic-stealth performance indicators (such as L/D , RCS). This requires extending the diffusion model to a conditional denoising diffusion probabilistic model (CDDPM).

The core idea of condition embedding is to use the condition information c' (i.e., the target performance indicator vector) as an additional input to guide the entire diffusion process (including training and sampling). The noise prediction network is thus extended to $\epsilon_\theta(x_t, t, c')$. The training loss function is modified accordingly as

$$\mathcal{L}_{\text{cond}} = E_{x_0, \epsilon \sim N(0, I), t} [\|\epsilon - \epsilon_\theta(x_t, t, c')\|^2] \quad (13)$$

While learning to denoise, the network must also learn how condition c' influences data generation, thereby establishing a mapping from the performance space to the design space.

Network architecture The noise prediction network ϵ_θ serves as the core component of the CDDPM. In this work, a U-Net-based architecture^[32] is adopted. The model follows an encoder-decoder structure: The encoder captures contextual information through down-sampling, while the decoder reconstructs spatial details via up-sampling and skip connections. This design is well-suited for processing spatially structured data such as images or sequences of CST parameters representing airfoils.

The condition vector c , comprising target performance parameters (L/D and RCS), is first transformed into a feature embedding via a multilayer perceptron (MLP). This embedding is then injected at multiple levels of the U-Net—typically through feature addition or concatenation—to ensure conditional guidance throughout the denoising process. The diffusion step t is incorporated into the network

using sinusoidal positional encoding.

The input to the network is processed by the encoder, which consists of four down-sampling stages based on residual convolutional blocks. These progressively reduce the spatial length (256 to 32) while increasing the number of feature channels (128 to 256). At the bottleneck, the representation is compressed to a length of 32 with 256 channels.

The decoder mirrors this structure through four up-sampling stages (32 to 256), implemented with transposed convolutions. To preserve fine-grained details, skip connections are employed between corresponding encoder and decoder layers, enabling the fusion of high-level semantic information with low-level spatial features and thus alleviating information loss during compression.

For condition injection, a two-layer MLP encodes the target performance parameters into an embedding vector c_{emb} , which is combined with the time embedding t_{emb} and added to each residual block. This ensures that the generation process remains consistently guided by the specified aerodynamic and stealth performance objectives.

This architecture ensures that both aerodynamic and stealth performance targets guide the generation process at every resolution level. By fusing condition embeddings with time embeddings, the network learns to associate CST parameters with target L/D and RCS values, enabling the direct performance-to-shape mapping that distinguishes our approach from conventional surrogate-based methods.

1.3 Training method

In this study, the diffusion process is set to 1 000 steps with a linear noise schedule where $\beta_{\text{start}} = 10^{-4}$ and $\beta_{\text{end}} = 0.02$. Time step information is encoded using a sinusoidal time embedding layer. Condition information—comprising target L/D and RCS values—is projected into a high-dimensional feature space via a condition projection layer, enabling effective multimodal information fusion.

The backbone noise prediction network is structured as a 6-layer fully connected network, with 256 neurons in each layer, and uses the SiLU activation

function. The time embedding is implemented using sine-cosine positional encoding with a dimension of 256. The condition projection layer maps the 2-D target performance vector to a 256-D feature representation.

Training is performed using the AdamW optimizer with a learning rate of 10^{-4} and weight decay of 10^{-6} . The loss function is defined as MSE between the predicted and actual noise. The model is trained with a batch size of 64 for a total of 20 000 epochs. The training loss of the CDDPM in Fig.6 converges relatively slowly, reflecting its need for more time to learn the complex data distribution and conditional mapping relationships. However, this also implies that it possesses stronger design space exploration capability during the generation process, which is beneficial for generating airfoils that meet complex design constraints.




Fig.6 Training loss convergence curve of the CDDPM

2 Results of CDDPM-Based Generative Design

2.1 Generation results of CDDPM

2.1.1 Evaluation cases design

To systematically evaluate the performance of the CDDPM in airfoils generation, five typical aerodynamic-stealth performance combinations are selected as evaluation benchmarks. Each group contains six independent design cases, resulting in a total of 30 test samples. The selection of test cases follows three main principles.

(1) Coverage of representative conditions: Four representative combinations of L/D and RCS are chosen, spanning high and low values of both

metrics. An additional challenging target with higher L/D and lower RCS is included as a breakthrough group, intended to explore performance beyond the original dataset. The five condition groups are specified in Table 1.

Table 1 Selected design cases for evaluation

Condition group	L/D	RCS/m ²
Group 1	4.71	0.61
Group 2	4.71	0.07
Group 3	14.63	0.61
Group 4	14.63	0.07
Group 5	15.07	0.068

This design ensures that the evaluation covers typical trade-off scenarios between aerodynamic efficiency and stealth performance commonly encountered in practical engineering applications.

(2) Verification of statistical robustness: Each performance group contains six independently generated airfoils. By evaluating multiple design realizations under the same performance target, the model's stability and consistency under fixed constraints can be assessed, thereby mitigating the impact of random variations that may arise in single-sample evaluations.

(3) Investigation of multidisciplinary trade-offs: Particular emphasis is placed on analyzing the inherent trade-off between aerodynamic and stealth performance. By comparing design outcomes across different condition groups, the model's ability to handle coupled multidisciplinary constraints can be thoroughly examined.

2.1.2 Generation results and analysis

Based on the five evaluation condition groups, the target performance specifications were input into the trained CDDPM to generate corresponding airfoil designs. The resulting airfoils are presented in Fig.7. It can be observed that airfoils within the same group exhibit similar shape distributions while maintaining diversity in geometric features such as leading-edge radius, trailing-edge configuration, and maximum thickness location. The generated results reveal clear physical insights: Airfoils de-

signed for low RCS tend to feature sharper leading edges, whereas those optimized for high L/D display more rounded leading-edge contours and gentler lower-surface curvature. This morphological distinction aligns consistently with the inherent trade-off between aerodynamic efficiency and stealth performance, demonstrating the model's ability to capture underlying multidisciplinary design principles.




Fig.7 Generated airfoils for groups 1—5

Following geometric inspections that confirmed the smoothness and physical reasonableness of all generated airfoils, these designs were subsequently imported into VSPAero and FEKO for aerodynamic and stealth performance simulations. The actual L/D and RCS values obtained from these simulations were then compared against the preset target values for verification. To quantitatively evaluate the performance accuracy, both the maximum relative error and average error metrics were employed to assess the deviations between the achieved aerodynamic-stealth performance and the target specifications. Fig.8 presents the average relative errors between the simulated performance values and the target indicators across the five groups




Fig.8 Average relative errors of L/D and RCS for five groups

of designed airfoils.

Overall, the CDDPM demonstrated strong predictive accuracy across all 30 evaluation cases. Error analysis reveals that for the L/D , the maximum average relative error reached 3.14%, while the overall average relative error stood at 2.18%. For the RCS, the corresponding values were 5.78% and 3.92%, respectively. As illustrated in Fig.8, the average relative errors between the achieved aerodynamic-stealth performance and the target values remain generally small across all groups. Although the RCS errors tend to be slightly larger than those for L/D , all deviations are effectively controlled within a 6% threshold, confirming the model's robustness in meeting multidisciplinary design requirements.

To further evaluate the proposed CDDPM, we compare its inverse design accuracy against three deterministic regression models: Gaussian process regression (GPR) with a radial basis function (RBF) kernel, support vector regression (SVR) with an RBF kernel, and random forest (RF) with 100 trees. All models are trained on the same 1 360 samples to map target performance (L/D and RCS) to the 19 CST parameters. Unlike the generative CDDPM, which can produce multiple candidates per condition, these regression models output a single deterministic prediction. Accordingly, for the five representative condition groups defined in Section 2.1.1, each baseline model generates exactly five airfoils (one per group). These airfoils are then simulated using the same VLM and PO solvers to obtain their actual L/D and RCS, and the relative errors are calculated. The average relative errors across the five test cases are summarized in Table 2.

Table 2 Comparison of prediction accuracy between baseline models and CDDPM

Model	Average relative error/%	
	L/D	RCS
GPR	2.95	5.08
SVR	4.12	6.73
RF	3.58	6.21
CDDPM(ours)	2.18	3.92

As shown in Table 2, the CDDPM achieves the lowest average relative errors for both L/D and RCS, outperforming all baseline models. These results demonstrate that the proposed diffusion-based generative model provides superior predictive fidelity in inverse design under multidisciplinary constraints, attributed to its ability to model complex conditional distributions.

The observed phenomenon of RCS errors generally exceeding L/D errors primarily stems from inherent differences in their geometric sensitivity. As a core aerodynamic performance indicator, L/D represents a globally integrated quantity that is sensitive to overall airfoil characteristics such as camber and thickness distribution, with calculation results demonstrating high numerical stability. In contrast, RCS is a physical quantity highly sensitive to local geometric features—particularly leading-edge curvature and trailing-edge sharpness—with an electromagnetic response mechanism characterized by stronger nonlinearity. The CST parameterization method employed in this study effectively captures macroscopic geometric features governing aerodynamic performance but exhibits inherent limitations in controlling the fine local details that predominantly determine RCS values. Furthermore, uncertainties inherent in electromagnetic numerical simulations may introduce additional computational noise. These factors collectively render the “geometry-RCS” mapping relationship more complex than the “geometry- L/D ” correlation, consequently manifesting as slightly larger errors in the generative design process.

The results demonstrate that the CDDPM maintains high design accuracy across varying performance requirements, confirming its effectiveness and applicability in aerodynamic-stealth generative design tasks. Notably, even under the most challenging performance combination of higher L/D and lower RCS, the model preserves satisfactory design accuracy, highlighting its distinctive advantage in resolving multidisciplinary trade-off problems.

2.2 Optimization design results and discussion

2.2.1 Design space expansion via conditional generation

Building upon the trained and converged CD-DPM described in Section 1, a total of 2 000 wing airfoil CST parameter vectors were generated within a specified target performance space. The generation process was configured as follows: The standard DDPM sampling algorithm was employed, maintaining the 1 000-step sampling procedure consistent with the original training setup. The target performance space was designed to encompass the extreme values of the training set while allowing for moderate extrapolation, with $L/D \in [4.7, 15.7]$, $RCS \in [0.03, 0.63]$. This setup enabled the generation of samples both within and beyond the original dataset’s performance distribution, also ensured comprehensive exploration of the design space boundaries. All conditioning inputs were normalized before being fed into the diffusion model, and diverse airfoil CST parameters satisfying the target performance constraints were obtained through the complete inverse denoising process.

To assess the design space exploration capability of the generated samples, a performance space distribution comparison is presented in Fig.9. The main observations are summarized as follows.

(1) Original dataset (1 360 samples): Represented by blue scatter points in Fig.9, the original samples exhibit a moderately dispersed yet uneven distribution. Sparse sampling is observed in the low L/D region (comprising only 21.1% of samples),




Fig.9 Performance space comparison of the generated samples and the original ones

while exploration of the high L/D region remains inadequate. Most samples are concentrated in medium L/D and medium RCS regions, reflecting the limitations of traditional sampling methods in addressing multi-objective trade-off areas.

(2) Generated samples (2 000 samples) : Shown as red scatter points in Fig.9, the generated samples demonstrate improved global coverage. They uniformly cover the original data's performance distribution region, with a core density overlap exceeding 0.85. The Jensen-Shannon divergence (JSD) value of 0.143 indicates that the generated distribution remains reasonably consistent with the original. Furthermore, the model exhibits significant boundary-breaking capability, producing dense samples (accounting for 22.5%) in the high L/D -low RCS conflict region ($L/D > 14.63$ or $RCS < 0.07$), substantially extending beyond original data boundaries. The performance space coverage of generated samples is 1.32 times greater than that of the original dataset. At the same time, compared with the original dataset, the minimum RCS is reduced by 28.6%, and the maximum L/D is increased by 7.5%, with dense sampling achieved both within and beyond the original performance space.

The generation of 2 000 new airfoil samples using the trained CDDPM was completed in 12.3 minutes on a single NVIDIA V100 GPU (32 GB VRAM), corresponding to an average generation time of 0.37 s per airfoil. This efficiency highlights the potential of the proposed method for real-time

design space exploration, as the generation cost is negligible compared to the 22.7 core-hours required for the initial dataset simulation.

The distribution consistency test for the CST parameters of the generated samples is shown in Table 3.

Table 3 Comparison of CST parameter ranges between original and generated samples

Parameter	Original range	Generated range	Overlap
1	[0.073, 0.219]	[0.053, 0.221]	0.887 544
2	[0.062, 0.185]	[0.049, 0.194]	0.885 305
3	[0.080, 0.239]	[0.049, 0.247]	0.884 127
4	[0.043, 0.130]	[0.040, 0.131]	0.805 303
5	[0.088, 0.264]	[0.069, 0.283]	0.900 911
6	[0.039, 0.118]	[0.030, 0.123]	0.851 316
7	[0.074, 0.220]	[0.066, 0.252]	0.880 809
8	[0.053, 0.158]	[0.029, 0.163]	0.914 106
9	[0.060, 0.180]	[0.053, 0.191]	0.875 429
10	[0.058, 0.174]	[0.046, 0.194]	0.816 990
11	[−0.062, −0.185]	[−0.059, −0.193]	0.883 118
12	[−0.080, −0.239]	[−0.070, −0.255]	0.825 461
13	[−0.043, −0.130]	[−0.039, −0.138]	0.804 747
14	[−0.088, −0.264]	[−0.068, −0.293]	0.878 669
15	[−0.039, −0.118]	[−0.023, −0.123]	0.878 743
16	[−0.074, −0.220]	[−0.051, −0.231]	0.903 530
17	[−0.053, −0.158]	[−0.052, −0.163]	0.833 116
18	[−0.060, −0.180]	[−0.057, −0.195]	0.884 890
19	[−0.058, −0.174]	[−0.032, −0.180]	0.806 373
Average			0.862 657

The overlap ratios between the range of generated CST parameters and the range of original parameters all exceed 0.85, demonstrating that the generated distribution maintains high consistency with the original data. Compared to the original dataset, the parameter ranges have been significantly expanded, with substantial exploration achieved at both upper and lower parameter limits. Notably, the model successfully generated samples at the extremes of L/D and RCS within parameter ranges that maintain an average matching degree exceeding 86%.

As illustrated in Fig.10, a representative subset of 25 generated airfoils is presented. The L/D

and RCS values of these airfoils essentially cover the entire performance range, with all sampled points located in previously unrepresented regions of the original sample space.




Fig.10 Generated airfoils falling in the performance range

2.2.2 Evaluation of optimization design capability

To further evaluate the generative model's capability in aerodynamic-stealth optimization of wing airfoils, this study investigates whether the model can generate samples that dominate the existing pareto front, thereby achieving genuine performance improvement.

Analysis of the 2 000 generated samples reveals that the CDDPM successfully produced 339 samples that lie beyond the original Pareto front, with their distribution visualized in Fig.11. These superior samples include not only designs with significantly reduced RCS but also a considerable number of airfoils with enhanced L/D . Particularly noteworthy is the concentration of improved samples in the high L/D and low RCS boundary region, highlighting the model's ability to advance the performance frontier.




Fig.11 Distribution of generated samples relative to the original dataset Pareto front

Analysis of the 2 000 generated samples reveals that the CDDPM successfully produced 339 samples that lie beyond the original Pareto front, with their distribution visualized in Fig.11. These superior samples include not only designs with significantly reduced RCS but also a considerable number of airfoils with enhanced L/D . Particularly noteworthy is the concentration of improved samples in the high L/D and low RCS boundary region, highlighting the model's ability to advance the performance frontier.

When applying a strict optimization target of $L/D > 13.6$ and $RCS < 0.1 \text{ m}^2$ simultaneously, the model generated 22 high-performance airfoils—nearly triple the eight samples meeting these criteria in the original dataset. These findings demonstrate that the diffusion model possesses dual advantages in design sample generation: It not only approximates preset optimization targets with high accuracy but also maintains exceptional generation efficiency. This superior performance stems from the model's profound design space exploration capability during generation, enabling it to effectively capture both diversity and superiority characteristics within complex constraint spaces. The results align with the inherent advantages of diffusion models in sample diversity and innovative generation, fully validating their practical efficacy in addressing complex multidisciplinary optimization challenges.

2.2.3 Evaluation of optimization convergence

To quantitatively evaluate the convergence behavior of the proposed generative optimization framework, the hypervolume indicator^[33] is employed. Hypervolume measures the volume of the objective space dominated by a set of non-dominated solutions relative to a reference point, providing a comprehensive metric for both convergence and diversity of the Pareto front.

In this study, the reference point for hypervolume calculation is set as $[RCS_{\max}, (L/D)_{\min}] = [0.61, 4.71]$, which is determined by the worst feasible values observed in the original dataset to en-

sure a fair comparison. The hypervolume is computed for the non-dominated solutions obtained from both the original dataset and the CDDPM-generated sample set.

As illustrated in Fig.12, the hypervolume of the original dataset is calculated to be $HV_0=5.07$. After generating 2 000 new samples via the CDDPM, the hypervolume of the expanded Pareto front increases to $HV_1=5.91$, representing a relative improvement of 16.6%. This significant growth indicates that the generative model effectively drives the optimization process toward a more converged and expansive Pareto front.




Fig.12 Hypervolume evolution during the generative optimization process

Furthermore, to demonstrate the iterative convergence behavior, we monitor the hypervolume growth as a function of the number of generated samples. The results show that the hypervolume stabilizes after generating approximately 1 200 samples, suggesting that the diffusion model has sufficiently explored the feasible design space and converged to a stable Pareto frontier. This convergence analysis confirms that the CDDPM not only enhances the design space but also ensures reliable and reproducible optimization outcomes.

3 Conclusions

This study has developed an end-to-end framework for aerodynamic-stealth generative design of wing airfoils based on CDDPM. The CDDPM efficiently and accurately achieves diversified mapping from target aerodynamic performance (L/D) and stealth performance (RCS) to airfoil geometric pa-

rameters (CST parameters). The tested results show that the average relative errors for L/D and RCS were smaller than 2.18% and 3.92%, respectively, demonstrating the model's powerful capability in handling high-dimensional nonlinear inverse design problems.

In the CDDPM, the joint aerodynamic-stealth conditional injection mechanism effectively coordinates the performance conflicts between L/D and RCS. Test results indicate that the model maintains high design consistency and stability under various performance constraints. Notably, even under the most challenging high L/D (15.07) and low RCS (0.068 m^2) conditions, the average relative errors for L/D and RCS are controlled within 3.14% and 5.78%, respectively, underscoring its robustness in multidisciplinary trade-off scenarios.

The CDDPM exhibits excellent design space expansion capabilities. By generating 2 000 new airfoil samples conditioned on a target performance space that extends beyond the original dataset, the model produces design that not only replicate the original statistical distribution (Jensen-Shannon divergence of 0.143) but also significantly push the performance boundaries. Compared to the original dataset of 1 360 samples, the generated samples enlarge the design space coverage by a factor of 1.32, achieving a minimum RCS of 0.068 m^2 (a 28.6% reduction) and a maximum L/D of 15.07 (a 7.5% increase). Importantly, 339 of the generated samples (17%) lie beyond the original Pareto front, demonstrating the model's ability to discover superior designs. For the stringent simultaneous requirement of $L/D > 13.6$ and $\text{RCS} < 0.1 \text{ m}^2$, the model yields 22 high-performance airfoils, nearly tripling the eight samples meeting these criteria in the original dataset. The hypervolume indicator, measuring the dominated objective space, improves from 5.07 (original) to 5.91 (generated), a relative gain of 16.6%, confirming the enhanced convergence and diversity of the expanded Pareto front.

This research establishes a workflow from au-

tomated modeling and simulation to generative design. It provides a reliable and efficient new methodology for aerodynamic-stealth integrated airfoil design. The current study focuses on wing airfoils. In the future, the CDDPM can be applied to three-dimensional aircraft generative design. This expansion faces the challenges such as higher-dimensional parameterization, escalating computational costs, and more complex fluid-electromagnetic coupling effects. Furthermore, constraints including structural integrity, flight stability, and manufacturability should be considered for aircraft design.

References

- [1] ALIZADEH R, ALLEN J K, MISTREE F. Managing computational complexity using surrogate models: a critical review[J]. *Research in Engineering Design*, 2020, 31(3): 275-298.
- [2] SHARMA S, KUMAR V. A comprehensive review on multi-objective optimization techniques: Past, present and future[J]. *Archives of Computational Methods in Engineering*, 2022, 29(7): 5605-5633.
- [3] GONG J D, LIU Q W, WANG K. Gradient descent algorithm for small UAV parameter estimation system [J]. *Transactions of Nanjing University of Aeronautics and Astronautics*, 2017, 34(6): 680-687.
- [4] SHAN S, WANG G G. Survey of modeling and optimization strategies to solve high-dimensional design problems with computationally-expensive black-box functions[J]. *Structural and Multidisciplinary Optimization*, 2010, 41(2): 219-241.
- [5] BACH F. Breaking the curse of dimensionality with convex neural networks[J]. *Journal of Machine Learning Research*, 2017, 18(19): 1-53.
- [6] WU H Z, LIU X J, AN W, et al. A generative deep learning framework for airfoil flow field prediction with sparse data[J]. *Chinese Journal of Aeronautics*, 2022, 35(1): 470-484.
- [7] CHOWDHARY K R. Natural language processing [M]. New Delhi: Springer, 2020: 603-649.
- [8] SZELISKI R. Computer vision: Algorithms and applications[M]. 2nd ed. Cham: Springer Nature, 2022.
- [9] LEADER J J. Numerical analysis and scientific computation[M]. Boca Raton: Chapman and Hall/CRC, 2022.
- [10] NIU J W, LIU Z G, PAN Q, et al. Conditional self-attention generative adversarial network with differential evolution algorithm for imbalanced data classification[J]. *Chinese Journal of Aeronautics*, 2023, 36(3): 303-315.
- [11] CAO H Q, TAN C, GAO Z Y, et al. A survey on generative diffusion models[J]. *IEEE Transactions on Knowledge and Data Engineering*, 2024, 36(7): 2814-2830.
- [12] WU Z J, GE Y X, WANG X T, et al. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation[C]//*Proceedings of the IEEE/CVF International Conference on Computer Vision*. Paris, France: IEEE, 2023: 7623-7633.
- [13] HO J, JAIN A, ABBEEL P. Denoising diffusion probabilistic models[J]. *Advances in Neural Information Processing Systems*, 2020, 33: 6840-6851.
- [14] AGGARWAL A, MITTAL M, BATTINENI G. Generative adversarial network: An overview of theory and applications[J]. *International Journal of Information Management Data Insights*, 2021, 1(1): 100004.
- [15] WANG Y, BLEI D, CUNNINGHAM J P. Posterior collapse and latent variable non-identifiability[J]. *Advances in Neural Information Processing Systems*, 2021, 34: 5443-5455.
- [16] MAZÉ F, AHMED F. Diffusion models beat GANs on topology optimization[C]//*Proceedings of the AAAI Conference on Artificial Intelligence*. Washington, DC: AAAI Press, 2023, 37(8): 9108-9116.
- [17] GONG Zhiyuan, ZHAO Huan. Latent diffusion model for multidisciplinary and multi-objective airfoil inverse design[EB/OL]. (2024-08-26) [2025-11-01]. <https://doi.org/10.36227/techrxiv.172468827.77833794/v1>.
- [18] WANG Jing, LIU Wei, XIE Hairun, et al. Diffusion model-driven multi-objective generative design of supercritical airfoils[J]. *Acta Aeronautica et Astronautica Sinica*, 2025, 46(10): 155-167. (in Chinese)
- [19] CHENG M, HE L, LIN J Z, et al. Research on fast inverse design of airfoil based on denoising diffusion model[J]. *Journal of Aerospace Power*, 2026, 41(4): 1-10.
- [20] ZHANG Ruitao, WANG Cong, TAO Jun, et al. Airfoil parameterization method based on latent diffusion model[J]. *Acta Aeronautica et Astronautica Sinica*, 2025, 46(10): 111-125. (in Chinese)
- [21] TAO G C, LIU Y, YAN Y. Inverse design of transonic airfoils using denoising diffusion probabilistic

- models[C]//Proceedings of the 15th Asia Conference on Mechanical and Aerospace Engineering (ACMAE 2024). Harbin, China:[s.n.], 2025; 9-19.
- [22] GRAVES R, FARIMANI A B. Airfoil diffusion: Denoising diffusion model for conditional airfoil generation[EB/OL]. (2024-12-18) [2025-11-01]. <https://doi.org/10.48550/arXiv.2408.15898>.
- [23] YONEKURA K, OSHIMA Y, AICHI M. Preliminary study of airfoil design synthesis using a conditional diffusion model and smoothing method[J]. *Computation*, 2024, 12(11): 227.
- [24] XU Y S, TONG X H, STILLA U. Voxel-based representation of 3D point clouds: Methods, applications, and its potential use in the construction industry [J]. *Automation in Construction*, 2021, 126: 103675.
- [25] NALLAPATI R. Discriminative models for information retrieval[C]//Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval. New York: ACM, 2004: 64-71.
- [26] AKRAM M T, KIM M H. CFD analysis and shape optimization of airfoils using class shape transformation and genetic algorithm—Part I [J]. *Applied Sciences*, 2021, 11(9): 3791.
- [27] HELTON J C, DAVIS F J. Latin hypercube sampling and the propagation of uncertainty in analyses of complex systems[J]. *Reliability Engineering & System Safety*, 2003, 81(1): 23-69.
- [28] MIMEAU C, MORTAZAVI I. A review of vortex methods and their applications: From creation to recent advances[J]. *Fluids*, 2021, 6(2): 68.
- [29] LIU Z L, WANG C F. Efficient iterative method of moments—Physical optics hybrid technique for electrically large objects[J]. *IEEE Transactions on Antennas and Propagation*, 2012, 60(7): 3520-3525.
- [30] JR ANDERSON J D, CORDA S, VAN WIE D M. Numerical lifting line theory applied to drooped leading-edge wings below and above stall[J]. *Journal of Aircraft*, 1980, 17(12): 898-904.
- [31] DORF R C. The electrical engineering handbook[M]. Boston: Elsevier Academic Press, 1997.
- [32] RONNEBERGER O, FISCHER P, BROX T. U-net: Convolutional networks for biomedical image segmentation[C]//Proceedings of the International Conference on Medical Image Computing and Computer-Assisted Intervention. Cham: Springer International Publishing, 2015: 234-241.
- [33] ZITZLER E, THIELE L. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach[J]. *IEEE Transactions on Evolutionary Computation*, 1999, 3(4): 257-271.

Acknowledgements This work was supported by the Fundamental Research Funds for the Central Universities (No. 226-2025-00075), the Defense Industrial Technology Development Program (No.JCKY2023205B013); and the Leading Talent Project for Scientific and Technological Innovation in Zhejiang Province (No.2023R5220).

Authors

The first author Mr. CHI Xinyan received his B.S. degree in engineering from Zhejiang University in 2024 and is currently a M.S. student at School of Aeronautics and Astronautics of Zhejiang University. His main research direction is intelligent design of aircraft.

The corresponding author Dr. ZENG Lifang received the B.S. degree in aircraft design and engineering and the dual M.S. degrees in aircraft design and engineering management from Nanjing University of Aeronautics and Astronautics (NUAA) and Hamburg University of Technology (TUHH), and the Ph.D. degree in fluid mechanics from Zhejiang University. She is currently an assistant researcher with the Institute of Intelligent Unmanned Systems, School of Aeronautics and Astronautics, Zhejiang University. Her current research focuses on the aerodynamics of rotating bodies and bio-inspired flight, and she is currently hosting a project supported by the National Natural Science Foundation of China.

Author contributions Mr. CHI Xinyan designed the study, compiled the models, conducted the analysis, interpreted the results, and wrote the manuscript. Dr. ZENG Lifang provided a validation of the study and made revisions to the manuscript. Prof. LI Jun contributed to the discussion and background of the study. Mr. LI Yuhang provided technical support. All authors commented on the manuscript draft and approved the submission.

Competing interests The authors declare no competing interests.

基于生成式模型的机翼翼型气动隐身一体化优化设计方法

池心砚¹, 曾丽芳^{1,2}, 黎 军^{1,3}, 李宇航¹

(1. 浙江大学航空航天学院, 杭州 310027, 中国; 2. 浙江大学浙江省智能高速无人机技术重点实验室, 杭州 310027, 中国; 3. 浣江实验室, 诸暨 311800, 中国)

摘要:生成式设计方法已被广泛应用于现代飞行器气动设计领域。然而当前气动-隐身一体化优化设计仍依赖代理模型和多目标优化算法。为了解决当前飞行器气动-隐身设计中繁杂的验证和优化流程,本文提出一种基于条件去噪扩散概率模型(Conditional denoising diffusion probability model, CDDPM)的快速生成式设计方法。首先采用类形状变化(Class-shape transformation, CST)方法对翼型进行参数化建模;基于涡格法(Vortex lattice method, VLM)和大尺寸物体物理光学(Physical optics, PO)法分别计算升阻比(L/D)和雷达散射截面(Radar cross-section, RCS),构建多学科性能数据集;基于数据集构建CDDPM,实现从目标性能到CST参数的端到端逆向映射。结果表明,CDDPM在气动隐身性能指标上的平均误差分别为2.18%和3.92%,所有误差控制在6%以内,同时生成的翼型也保有多样性。进一步地,生成样本将设计空间扩大1.32倍,最小RCS降低28.6%,最大 L/D 提升7.5%。本文提出的一种基于生成式模型的翼型气动隐身快速优化设计方法框架,为AI驱动飞行器多学科设计优化提供了新途径。

关键词:气动-隐身一体化设计;生成式设计;扩散模型;类形状变化参数化;升阻比;雷达散射截面

研究亮点:

1. 将条件去噪扩散概率模型应用于翼型气动-隐身一体化优化设计,实现了从目标性能到翼型几何参数的端到端逆向映射。
2. 提出联合条件注入机制,将气动与隐身双重目标同时嵌入扩散过程,通过生成式模型主动扩展设计空间,有效协调了升阻比与雷达散射截面之间的内在冲突。