In: Proceedings of the 32nd AIAA/CEAS Aeroacoustics Conference (2026), 26-29 May 2026, Brussels, Belgium, Article number 2026-3451
DOI: 10.2514/6.2026-3451

Machine Learning Approaches for Preliminary Design of Rotorcraft Aeroacoustics

Peekarn Teerawathananon*, Ioannis Goulos.† and David MacManus‡
Cranfield University, Cranfield, United Kingdom, MK43 0AL

Rotorcraft noise prediction remains challenging due to the large cost of experimental campaigns and high-order aeroacoustic simulations. This work develops a reduced-order, deep learning-based framework for the rapid synthesis of helicopter main-rotor noise hemispheres for preliminary design. A rotorcraft aero-mechanics and acoustics computational framework was used to generate noise hemispheres based on Latin Hypercube Sampling (LHS) across a wide range of design parameters and operating conditions. Two reduced-order modelling (ROM) methods were investigated: (i) POD + Kriging, combining dimensionality-reduction with Gaussian-process regression, and (ii) an enhanced U-Net Convolutional Neural Network (U-CNN) that reconstructs the full acoustic field directly from rotor design parameters. The ROMs were validated using cross-validation and evaluated on low-speed descent cases, a condition typically dominated by complex Blade-Vortex Interaction (BVI) noise. The POD-Kriging model achieves a mean root-mean-square error of 0.70 dB and the U-CNN achieves 0.53 dB across the unseen test configurations, recovering thickness and loading noise components even in BVI-dominated cases while maintaining the original resolution as a traditional computation method. Overall, the study establishes a novel framework for the rapid rotor noise characterisation suitable for rotorcraft preliminary design.

I. Nomenclature

Roman Symbols

c_b	=	Blade chord, m	M_{tip}	=	Blade-tip Mach number
C_L	=	Rotor disc rolling moment coefficient	N_b	=	Number of blades
C_M	=	Rotor disc pitching moment coefficient	N_s	=	Number of database samples
C_T	=	Rotor disc thrust coefficient	R	=	Main rotor radius, m
C_T/σ	=	Blade loading	R^2	=	Coefficient of determination
V_∞	=	Freestream velocity, m/s			

Greek Symbols

α_k	=	k-th POD modal coefficient	σ	=	Rotor solidity
β	=	Acoustic hemisphere elevation angle	Φ	=	POD basis (spatial modes)
θ_{tw}	=	Blade pre-twist angle, deg	ψ	=	Acoustic hemisphere azimuth angle
λ_{taper}	=	Planform taper ratio	Ω	=	Rotor rotational speed, rad/s
μ	=	Rotor advance ratio			

Acronyms

BVI	=	Blade-Vortex Interaction	POD	=	Proper Orthogonal Decomposition
MaxAE	=	Maximum Absolute Error, dB	RMSE	=	Root-Mean-Square Error, dB
OASPL	=	Overall Sound Pressure Level, dB	ROM	=	Reduced-Order Model
U-CNN	=	Enhanced U-Net CNN			

*PhD Researcher, Propulsion Engineering Centre, Cranfield University, p.teerawathananon@cranfield.ac.uk
†Senior Lecturer in Propulsion Aerodynamics and Performance Engineering, Propulsion Engineering Centre, Cranfield University
‡Head of Gas Turbine Technology Group, Centre for Propulsion and Thermal Power Engineering, Cranfield University

Published by AIAA. This is the Author Accepted Manuscript issued with: Creative Commons Attribution License (CC:BY 4.0).
The final published version (version of record) is available online at DOI:10.2514/6.2026-3451. Please refer to any applicable publisher terms of use.

II. Introduction

To meet the challenging performance and environmental targets defined by the Net Zero commitments [1] and civil noise-reduction roadmaps as Flightpath 2050 [2], accurate and efficient prediction of rotorcraft aeromechanic and aeroacoustic performance during the early stages of vehicle design is essential. The dominant contributors to helicopter noise are the main and tail rotors [3], and the resulting acoustic radiation is generally categorised into discrete-frequency (deterministic) and broadband (non-deterministic) components. Discrete-frequency noise arises from periodic aerodynamic phenomena and comprises thickness noise, loading noise, and high-speed impulsive (HSI) noise. Thickness noise is associated with the displacement of fluid due to the physical motion of the rotor blades, while loading noise results from unsteady aerodynamic forces exerted on the blade surfaces and contains both harmonic and impulsive contributions. The harmonic component corresponds to regular, periodic blade motion and is observed at multiples of the blade-passing frequency; the impulsive component becomes prominent during blade-vortex interaction (BVI) events, wherein a rotor blade intersects the tip vortex shed by a preceding blade, producing sharp, directional acoustic impulses. HSI noise occurs under transonic flow conditions at the advancing blade tip during high-speed flight and is characterised by localised shock formation and rapid pressure fluctuations, resulting in high-amplitude, short-duration acoustic pulses. The broadband component, which is aperiodic and stochastic in nature, arises from turbulent boundary-layer interaction with the blade surface (self-noise) and ingestion of atmospheric turbulence. The overall sound pressure level (OASPL) hemispheres predicted in the present work include discrete-frequency and broadband noise contributions; the modelling approach for each component is detailed in Section III.A. The acoustic characteristics of an isolated rotor are governed by numerous design and operational parameters, including blade count, rotor radius, rotational speed, blade chord, twist, taper, blade-tip thickness-to-chord ratio, and flight condition, motivating the need for models that can efficiently explore this multidimensional design space.

Current methods for the aeroacoustic characterisation of rotary-wing aircraft are based on the generation of acoustic hemispheres or ground noise footprints using either dedicated experiments (e.g. wind-tunnel testing with microphone arrays) [4] or high-fidelity simulations including Reynolds-Averaged Navier-Stokes (RANS) Computational Fluid Dynamics (CFD), Lattice Boltzmann (LB), and Large Eddy Simulation (LES) methods. These approaches are resource-intensive and generally reserved for the intermediate or detailed design stage, where the rotor geometry is already fixed [5], and are not suitable for tasks such as uncertainty quantification and design optimisation. Consequently, aeroacoustic performance is typically absent from the preliminary design trade space, increasing the risk of costly late-stage redesign when configurations fail to meet detectability requirements (for defence applications) or certification noise standards (for civil use). The challenge is further complicated by aerodynamic interactions between multi-rotor configurations. In helicopters, these interactions arise when the main rotor wake impinges on the tail rotor under specific operating conditions dictated by the vehicle's design characteristics [5]. Although the present work focuses on an isolated main rotor, the methodology is designed to be extensible to interactional noise in future work.

To accelerate acoustic characterisation, researchers have applied data-driven techniques to rotorcraft and propeller noise prediction. Existing studies have used artificial neural networks, regression models, and evolutionary algorithms to predict aggregate noise metrics, such as OASPL or spectral sound pressure levels at single or isolated observer locations, from rotor design parameters and operating conditions [6–8]; these approaches are computationally efficient but provide no information about how the noise is distributed spatially around the rotor. Surrogate models have also been developed to predict acoustic waveforms or tonal noise at discrete observer positions [9–11], recovering temporal or spectral detail but not the continuous directivity field between observers. For mission-level noise assessment, conventional ground footprint codes rely on pre-computed libraries of acoustic hemispheres corresponding to a finite set of flight states, with interpolation used to approximate intermediate conditions [12]; such libraries are inherently limited to the rotor configurations for which hemispheres have been generated. Physics-based approaches that compute hemispheres on-the-fly for each trajectory waypoint remove this restriction [13], but the computational cost of the underlying free-wake aeroelastic and acoustic solvers makes them impractical for the evaluations required during design-space exploration. Rotorcraft noise radiation is, however, inherently directional: BVI noise concentrates in narrow lobes at specific azimuth and elevation angles, thickness noise radiates predominantly in the rotor tip-path plane, and loading noise propagates out of plane. The complete directivity field is therefore essential for computing ground noise footprints, assessing compliance with certification or detectability requirements at prescribed observer positions, and optimising flight trajectories for noise abatement. No existing ROM combines full spatial resolution of the acoustic hemisphere with the ability to vary rotor design parameters.

The most closely related work is that of Kwon et al. [14], who constructed a POD-Kriging reduced-order model (ROM) to predict aerodynamic load distributions of an AH-1G rotor from operational parameters (advance ratio, descent angle, and control inputs), and then computed noise on the lower hemisphere using Farassat's formulation 1A [15].

This achieved average OASPL errors of 0.47 dB across 271 receivers in descent flight. Their framework enables rapid trajectory noise assessment and flight path optimisation for a fixed rotor geometry but does not support exploration of the rotor design space. The present work addresses this gap by developing and comparing two ROMs: (i) proper orthogonal decomposition with Kriging regression (POD-Kriging) and (ii) an enhanced U-Net convolutional neural network (U-CNN) that reconstruct the complete OASPL hemisphere directly from six rotor design variables, enabling acoustic assessment during the preliminary design phase when the rotor geometry is not yet fixed.

III. Methodology

A. Integrated Computational Framework

A computational framework for the aeroacoustic characterisation of helicopter rotors has been developed by integrating a series of validated tools for rotor aero-mechanics and acoustics developed at Cranfield University. The framework builds upon the HECTOR platform (Helicopter Omni-disciplinary Research Platform), originally developed by Goulos et al. [16–19] for mission-level performance evaluation of rotary-wing platforms. HECTOR encompasses flight dynamics modelling, engine performance prediction (via the Turbomatch code [20]), unsteady aerodynamics, aeroelastic modelling of rotor blades in flap-lag-torsion, and rotor noise prediction. For the present application, a nonlinear free-wake aeroelastic rotor model has been coupled with the ICARUS rotor noise prediction code (Integral Computational AeRoacoustics Unified Software) [13, 21]. The aeroelastic formulation integrates structural flexibility, unsteady blade aerodynamics, finite-state inflow, and dynamic wake curvature effects [22]. The free-wake rotor model

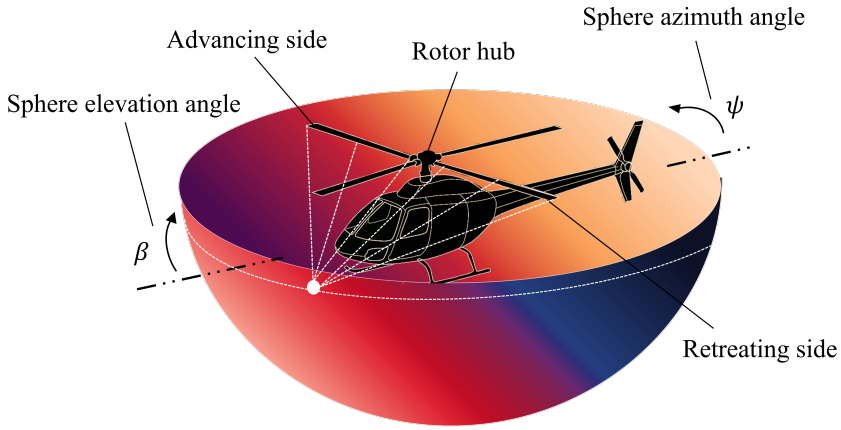


Fig. 1 Acoustic hemisphere definition.

provides high-resolution predictions of unsteady aerodynamic loads, capturing phenomena such as BVI and wake impingement. These loads are passed to the ICARUS solver, which computes the acoustic pressure using an integral formulation of the Ffowcs Williams-Hawkings (FW-H) equations [15, 23]. The formulation resolves the monopole (thickness) and dipole (loading) surface source terms. The quadrupole volume term, associated with HSI noise, is neglected under the assumption of low compressibility; this simplification is valid when blade-tip Mach numbers remain below the transonic regime, as is typical of the descent flight conditions and rotor configurations considered here. Broadband noise, which arises from stochastic sources such as boundary-layer self-noise, inflow turbulence interaction, and vortex shedding, requires detailed turbulence statistics for accurate first-principles prediction. This level of fidelity is computationally impractical when applied across the full design space. Instead, the ICARUS uses the semi-empirical model of Pegg [24] for the prediction of broadband noise, which is then energy-summed with the total contributions to obtain the OASPL. Previous investigations have shown that incorporating this model substantially improves agreement with ground-based acoustic measurements, particularly in regimes where neglecting broadband content leads to systematic underprediction [25]. The combined tonal and broadband contributions yield the OASPL on an acoustic hemisphere centred on the rotor hub. The hemisphere radius is set to 60 m, corresponding to approximately 15 rotor diameters for the model-scale rotor investigated (Sec. III.B), ensuring that far-field conditions are satisfied for the FW-H formulation [12]. The hemisphere is discretised into $N_\psi = 41$ azimuthal and $N_\beta = 26$ elevation stations, following the resolution adopted by Vouros [21], yielding a total of $N_\psi \times N_\beta = 1066$ observer points per configuration.

The OASPL is expressed as a function of azimuth (ψ) and elevation (β) on this spherical observer grid, as illustrated in Figure 1.

The free-wake aeroelastic and the ICARUS models has been validated against experimental measurements of unsteady airloads [26], aeroelastic blade deformations [22], and rotor noise predictions [13, 21] for the HART II rotor configuration. Further description of the individual models is available in the cited references; the focus here is on the numerical procedure and integration strategy.

B. Design of Experiments and Definition of Rotor Feature Space

The hingeless HART II rotor model [4] is used as the baseline rotor configuration. The HART II blade is a geometrically and dynamically scaled model of the BO-105 helicopter rotor, designed to match full-scale tip Mach number at 40% geometric scale. It was tested in the Higher Harmonic Control Aeroacoustics Rotor Test (HART II) programme [27] conducted in the German-Dutch Wind Tunnels (DNW), producing a comprehensive public dataset of blade pressures, structural loads, wake geometries, and acoustic measurements [26, 28, 29] There is a wide range of

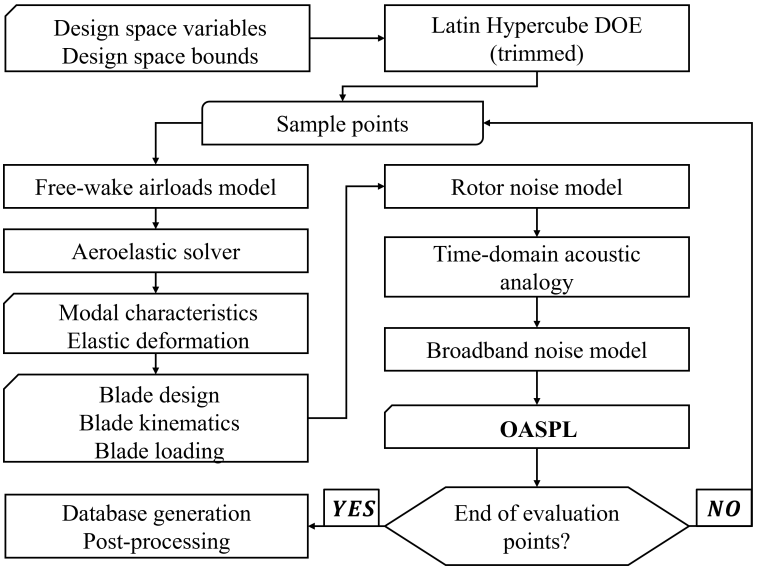


Fig. 2 Simulation procedure for trimmed rotor aeroacoustic prediction at each design point.

parameters influencing the tonal noise characteristics of a main rotor, including geometric features of the rotor blades and targeted operating conditions. Six rotor design parameters are varied in this study: the number of blades (N_b), rotor radius (R), blade chord (c_b), rotor speed (Ω), planform taper ratio (λ_{taper}), and blade pre-twist (θ_{tw}). Together, R and Ω set the blade-tip Mach number and thus govern thickness noise and the onset of compressibility effects; N_b controls the blade-passing frequency and the per-blade loading share; c_b determines the rotor solidity and the resulting trade-off between blade-section life coefficients and blade-volume displacement; and θ_{tw} and λ_{taper} redistribute the spanwise loading, directly affecting tip-vortex strength and BVI noise intensity. These parameters were selected because they exert a strong and coupled influence on both aerodynamic performance and the acoustic signature of an isolated main rotor. Their bounds, summarised in Table 2, span the range of values representative of conventional helicopter main rotor designs. A Latin Hypercube Sampling (LHS) method [30] was used to discretise this six-dimensional design space, selected for its space-filling properties which ensure adequate coverage with a manageable number of samples.

The six design variables are not all independent. For a fixed freestream velocity (V_∞), the advance ratio $\mu = V_\infty/\Omega R$ and blade-tip Mach number $M_{tip} = \Omega R/a$ are both determined by the choice of R and Ω, so Ω cannot be prescribed as an independent bound; its feasible range is determined implicitly by the μ and M_{tip} constraints for each sampled R. At the low advance ratios considered in this study, M_{tip} closely approximates the advancing-tip Mach number and therefore governs the onset of compressibility effects at the blade tip. Additional constraints on solidity $\sigma = N_b\bar{c}/(\pi R)$ and blade loading coefficient C_T/σ further restrict the feasible region in the remaining dimensions. Figure 3 shows the three-dimensional design space and its projection onto the R-Ω plane, identifying the infeasible regions excluded by these constraints.

Table 2 Rotor design parameters and their variation bounds. The HART-II column denotes the baseline values.

Feature		HART-II	Lower bound	Upper bound
Number of blades (−)	N_b	4	3	6
Rotor radius (m)	R	2.0	1.6	2.4
Blade chord (m)	c_b	0.121	0.103	0.139
Rotor speed (rad/s)*	Ω	109.12	84.56	143.88
Taper ratio (−)	λ_{taper}	1.0	0.6	1.0
Rotor blade twist (deg)	θ_{tw}	6.24	0	18

A total of 1500 candidate configurations that satisfy the feasible are retained and passed to the trimmed free-wake aeroelastic simulation loop (Fig. 2). Each is analysed at a low-speed descent flight condition, where BVI is known to be most pronounced [4], and trimmed to match the dimensional thrust, rolling moment, and pitching moment $(T_{req}, L_{req}, M_{req})$ of the baseline HART II case: thrust coefficient $C_T = 4.57 \times 10^{-3}$, rolling moment coefficient $C_L = 1.386 \times 10^{-5}$ and pitching moment coefficient $C_M = -1.386 \times 10^{-5}$, respectively. These trim targets are held constant so that changes in the predicted acoustic field can be attributed solely to variations in the rotor design parameters. Of the 1500 feasible configurations, 1287 converge to a trimmed equilibrium and form the final database.

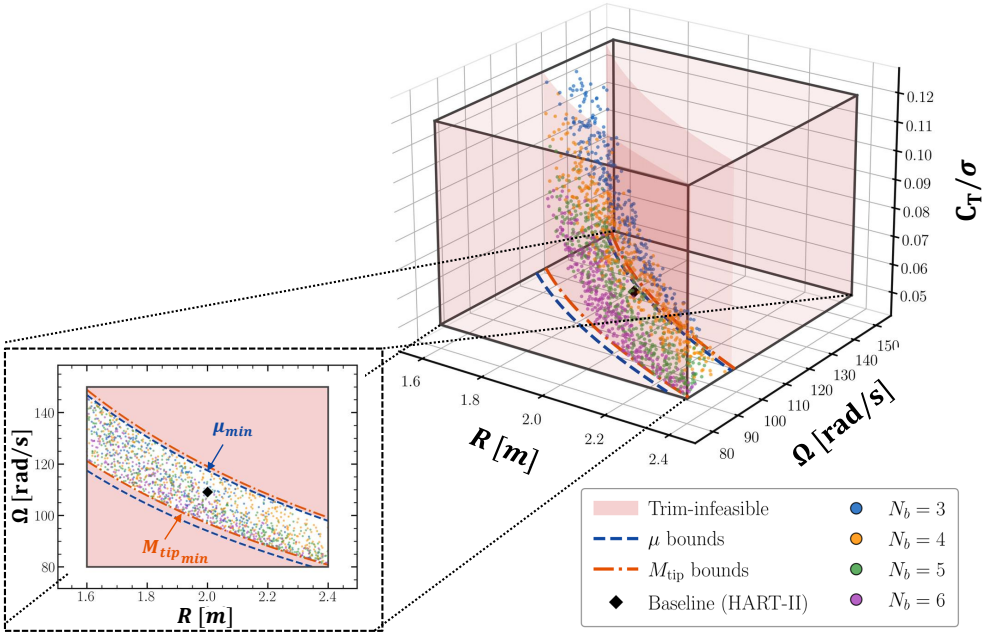


Fig. 3 Schematic showing the upper and lower bounds for the rotor design parameters (R, N_b, c_b) and operating conditions $(C_T/\sigma, \Omega, \mu, M_{tip})$

C. Proper Orthogonal Decomposition with Kriging Regression

The first reduced-order modelling approach combines dimensionality reduction through Proper Orthogonal Decomposition (POD) with Kriging regression to approximate the mapping from rotor design parameters to the acoustic hemisphere. The POD extracts a compact set of orthonormal spatial modes from the database, and a Kriging surrogate is then trained to predict the corresponding modal coefficients for new, untested rotor configurations. The two components are described in turn.

1. Proper orthogonal decomposition

The POD provides a systematic decomposition of the hemisphere database into a minimal set of energy-ranked orthonormal spatial modes, enabling a compact representation of the acoustic field [31, 32]. Each hemisphere in the database described in Section III.B represents the OASPL as a function of azimuth (ψ) and elevation (β) on the spherical observer grid (Fig. 1).

Let $\mathbf{p}^{(i)}(\psi, \beta) \in \mathbb{R}^{n_z}$ denote the vectorised OASPL field for the i^{th} sample. The complete dataset of n_s samples is assembled into a mean-subtracted snapshot matrix as

$$\mathbf{S} = \begin{bmatrix} \mathbf{p}^{(1)} - \overline{\mathbf{p}} \\ \mathbf{p}^{(2)} - \overline{\mathbf{p}} \\ \vdots \\ \mathbf{p}^{(n_s)} - \overline{\mathbf{p}} \end{bmatrix} \in \mathbb{R}^{n_s \times n_z}, \tag{1}$$

where $\overline{\mathbf{p}} = \frac{1}{n_s} \sum_{i=1}^{n_s} \mathbf{p}^{(i)}$ is the ensemble-mean noise field. Each centred sample is approximated as a linear combination of orthonormal spatial basis functions $\Phi_k(\psi, \beta)$, known as POD modes:

$$\mathbf{p}^{(i)}(\psi, \beta) \;\approx\; \overline{\mathbf{p}}(\psi, \beta) \;+\; \sum_{k=1}^{r} \alpha_k\big(\mathbf{x}^{(i)}\big)\, \Phi_k(\psi, \beta), \tag{2}$$

where $\alpha_k\big(\mathbf{x}^{(i)}\big)$ is the modal coefficient of mode k for the i-th sample, $r \ll n_z$ is the number of retained modes, and the input vector $\mathbf{x}^{(i)}$ consists of the six design variables defined in Table 2.

The decomposition is obtained by applying Singular Value Decomposition (SVD) to the snapshot matrix:

$$\mathbf{S} = \mathbf{U}\, \mathbf{\Sigma}\, \mathbf{V}^{\top}, \tag{3}$$

where $\mathbf{U} = [\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_{n_s}] \in \mathbb{R}^{n_s \times n_s}$ contains the left singular vectors, $\mathbf{V} = [\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_{n_z}] \in \mathbb{R}^{n_z \times n_z}$ contains the right singular vectors, and $\mathbf{\Sigma} \in \mathbb{R}^{n_s \times n_z}$ is a rectangular diagonal matrix with singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p \geq 0$ along its diagonal, where $p = \min(n_s, n_z)$. The POD spatial modes Φ_k correspond to the columns of $\mathbf{V}$ (right singular vectors), while the left singular vectors in $\mathbf{U}$ encode the temporal (parametric) variation across samples. The relation between the POD and SVD has been thoroughly established in [31, 32].

In practice, the economy-sized SVD is used, and the basis is truncated by retaining only the r most energetic modes such that a prescribed fraction η of the total energy is captured:

$$\frac{\sum_{k=1}^{r} \sigma_k^2}{\sum_{k=1}^{n_s} \sigma_k^2} \geq \eta. \tag{4}$$

The modal coefficients are obtained by projecting the mean-subtracted field onto the corresponding spatial mode:

$$\alpha_k\big(\mathbf{x}^{(i)}\big) = \mathbf{\Phi}_k^{\top}\, \big(\mathbf{p}^{(i)} - \overline{\mathbf{p}}\big), \tag{5}$$

or, equivalently, from the truncated SVD as $\mathbf{A} = \mathbf{U}_r\, \mathbf{\Sigma}_r$, where $\mathbf{U}_r \in \mathbb{R}^{n_s \times r}$ and $\mathbf{\Sigma}_r \in \mathbb{R}^{r \times r}$ are the truncated left-singular-vector and diagonal singular-value matrices, respectively. Each row of $\mathbf{A}$ corresponds to one design sample and contains its r modal amplitudes. The specific truncation level used in the present work is reported in Section IV.A.1.

2. Kriging regression

A Kriging surrogate model is constructed to approximate the nonlinear mapping between the rotor design parameters $\mathbf{x} \in \mathbb{R}^{n_{dv}}$ and each modal coefficient α_k obtained from the POD (Sec. III.C.1). The model is based on Gaussian Process Regression in the Universal Kriging formulation, with a quadratic trend function and an absolute exponential correlation kernel. The response $y(\mathbf{x})$ is expressed as the sum of a global trend function $f(\mathbf{x})$ and a zero-mean stochastic process $Z(\mathbf{x})$:

$$y(\mathbf{x}) = f(\mathbf{x}) + Z(\mathbf{x}). \tag{6}$$

The trend function is defined as a linear combination of second-order polynomial basis functions:

$$f(\mathbf{x}) = \sum_{l=1}^{p} \beta_l f_l(\mathbf{x}) = \mathbf{f}^T(\mathbf{x})\, \boldsymbol{\beta}, \tag{7}$$

where $\mathbf{f}(\mathbf{x}) \in \mathbb{R}^p$ is the vector of polynomial basis functions evaluated at $\mathbf{x}$, and $\boldsymbol{\beta} \in \mathbb{R}^p$ contains the regression coefficients.

The stochastic term $Z(\mathbf{x})$ captures local departures from the global trend. Its covariance is defined by [33]:

$$\mathrm{Cov}[Z(\mathbf{x}^i),\, Z(\mathbf{x}^j)] = \sigma^2\, R(\mathbf{x}^i,\, \mathbf{x}^j), \tag{8}$$

where σ^2 is the process variance and $R(\mathbf{x}^i,\, \mathbf{x}^j)$ is the correlation function between samples i and j. An absolute exponential correlation function is adopted following the default configuration of the Surrogate Modelling Toolbox [34]:

$$R(\mathbf{x}^i,\, \mathbf{x}^j) = \exp\left[-\sum_{d=1}^{n_{dv}} \theta_d\, |x_d^i - x_d^j|\right], \tag{9}$$

where $\theta_d > 0$ are hyperparameters controlling the correlation length along each design dimension d, and n_{dv} is the number of design variables. The $n_s \times n_s$ correlation matrix $\mathbf{R}$ is assembled with entries $R_{ij} = R(\mathbf{x}^i,\, \mathbf{x}^j)$. In practice, $\mathbf{R}$ is replaced by $\mathbf{R} + \delta\mathbf{I}$, where δ is a small nugget parameter that ensures numerical stability during factorisation; all subsequent expressions involving $\mathbf{R}$ include this regularisation.

The prediction at an untested design point $\mathbf{x}^*$ is given by:

$$\hat{y}(\mathbf{x}^*) = \mathbf{f}^T(\mathbf{x}^*)\,\hat{\boldsymbol{\beta}} + \mathbf{r}^T(\mathbf{x}^*)\,\mathbf{R}^{-1}\,(\mathbf{y} - \mathbf{F}\,\hat{\boldsymbol{\beta}}), \tag{10}$$

where $\mathbf{y} \in \mathbb{R}^{n_s}$ is the vector of training responses, $\mathbf{F}$ is the design matrix formed by evaluating the polynomial basis functions at the training points, and $\mathbf{r}(\mathbf{x}^*)$ is the correlation vector between the prediction point and all training points. The regression coefficients are obtained using the generalised least-squares estimator:

$$\hat{\boldsymbol{\beta}} = (\mathbf{F}^T\,\mathbf{R}^{-1}\,\mathbf{F})^{-1}\,\mathbf{F}^T\,\mathbf{R}^{-1}\,\mathbf{y}. \tag{11}$$

The hyperparameters $\boldsymbol{\theta}$ are estimated by maximising the concentrated log-likelihood function [33]:

$$\max(\boldsymbol{\theta}) = \; -\tfrac{1}{2}\left[n_s \ln(\hat{\sigma}^2) + \ln|\mathbf{R}|\right], \tag{12}$$

where both $\hat{\sigma}^2$ and $|\mathbf{R}|$ are functions of $\boldsymbol{\theta}$. The maximum-likelihood estimate of the process variance is given by

$$\hat{\sigma}^2 = \frac{1}{n_s}(\mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}})^T\mathbf{R}^{-1}(\mathbf{y} - \mathbf{F}\hat{\boldsymbol{\beta}}). \tag{13}$$

A separate Kriging model is trained for each of the r retained modal coefficients (Sec. III.C.1). For a new rotor configuration $\mathbf{x}^*$, the predicted modal coefficients are obtained from the r surrogates and substituted into the POD expansion (Eq. 2) to reconstruct the OASPL hemisphere:

$$\hat{\mathbf{p}}(\mathbf{x}^*;\psi,\beta) = \overline{\mathbf{p}}(\psi,\beta) + \sum_{k=1}^{r} \hat{\alpha}_k(\mathbf{x}^*)\,\Phi_k(\psi,\beta). \tag{14}$$

D. Enhanced U-Net convolutional neural networks

The second reduced-order model used a convolutional neural network (CNN) adapted from the U-Net architecture [35], originally developed for biomedical image segmentation, to predict the full OASPL hemisphere directly from rotor design parameters. Unlike the POD-Kriging method, this model predicts the entire two-dimensional directivity field in a single forward pass without requiring an intermediate decomposition. The U-Net is adopted as the backbone, with four enhancements incorporated to improve its suitability for structured acoustic field prediction: (i) Residual Dense Blocks (RDBs) [36] replace the standard convolutional blocks, (ii) Attention Gates (AG) [37] are applied to the skip connections, (iii) Convolutional Block Attention Modules (CBAM) [38] refine the decoder features, and (iv) circular padding enforces azimuthal periodicity. The resulting architecture is illustrated in Figure 4.

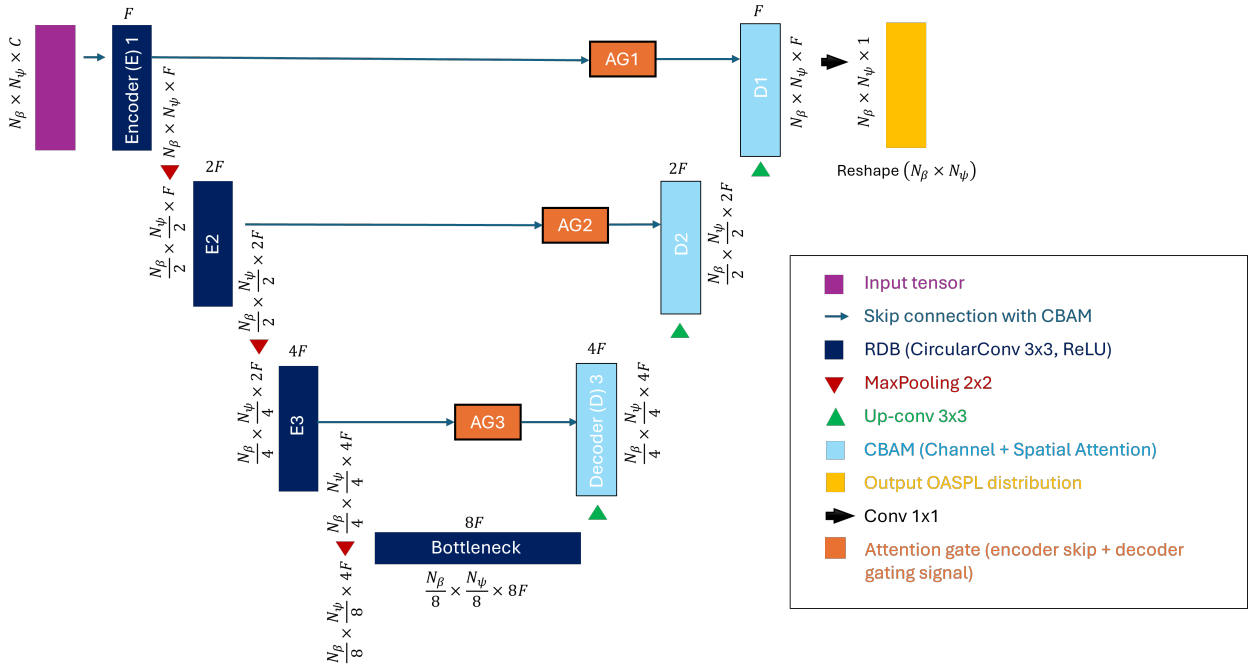


Fig. 4 Enhanced U-Net architecture for OASPL hemisphere prediction, adapted from [35].

Data preparation. The acoustic output from each simulation is a spatially structured field: the OASPL evaluated at N_β elevation angles and N_ψ azimuth angles, forming a grid of $N_\beta \times N_\psi$ observation points distributed over the hemisphere (Fig. 1). Neighbouring observation points exhibit strongly correlated acoustic levels due to the physically governed variation of rotor directivity. A standard multilayer perceptron (MLP) has no built-in mechanism to exploit these spatial correlations, treating the output as an unstructured vector [39]. A CNN, by contrast, operates directly on the two-dimensional grid using shared convolutional filters that capture local spatial dependencies by construction [40]. However, unlike CFD-based low order models that are typically trained on spatial meshes with inherent geometric information (e.g. Tejero et al. [41]), the present input consists solely of n_{dv} scalar design parameters (Tab. 2) with no spatial extent. To bridge this gap, the design vector is replicated at every point on the $N_\beta \times N_\psi$ grid and augmented with the azimuth and elevation coordinates of each grid point, producing an input tensor of shape $(N_\beta,\ N_\psi,\ C)$, where $C = n_{dv} + 2$ comprises n_{dv} channels for the design variables and two channels for the azimuth and elevation coordinates. The coordinate channels enable the network to associate each design parameter's influence with a specific location on the hemisphere; without them, the network cannot distinguish between observer locations and the output would be uniform across the hemisphere. The model output is a single-channel OASPL map of shape $(N_\beta,\ N_\psi)$. All input features and output fields are normalised to $[0,\ 1]$ using min-max scaling (Eq. 15).

$$\nu_{\text{norm}} = \frac{\nu - \nu_{\min}}{\nu_{\max} - \nu_{\min}} \tag{15}$$

where ν refers to any input or output variable.

Architecture. The model follows the symmetric encoder-decoder structure of the U-Net [35], in which skip connections between corresponding levels preserve spatial detail during reconstruction (Fig. 4). The encoder comprises three successive stages that extract features at progressively coarser spatial resolutions. At each stage, 2×2 max-pooling reduces the spatial dimensions by a factor of two, and the number of convolutional filters is doubled to compensate for the reduced spatial information [42]: the filter counts progress as F, $2F$, and $4F$ through the three encoder stages, with further doubling to $8F$ at the bottleneck, where $F = 32$ is the base filter count. All convolutional layers use 3×3 kernels with ReLU activation (Eq. 16). Feature extraction at each stage is performed by RDBs [36], in which each dense layer is extended with a learned sigmoid gate (Eq. 16) that enables the network to selectively weight individual feature

responses, attenuating activations that are not relevant to the prediction. Circular padding is applied along the azimuth axis to enforce continuity at $\psi = 0° \equiv 360°$, while reflective padding is used along the non-periodic elevation axis.

$$\phi_{\mathrm{ReLU}}(c_k) = \max(0,\, c_k), \qquad \phi_{\mathrm{Sigmoid}}(c_k) = \frac{1}{1 + e^{-c_k}}. \tag{16}$$

The decoder mirrors the encoder through three upsampling stages using 3×3 transposed convolutions with stride 2. At each level, Attention Gates [37] modulate the incoming skip connection based on the current decoder state, allowing the model to focus on spatially relevant encoder features. CBAM [38] further refine the decoder features through sequential channel and spatial attention, selectively weighting the most task-relevant feature channels and spatial regions. A final 1×1 convolution maps the refined features to the predicted OASPL field of shape $(N_\beta,\, N_\psi)$.

Training. The model is trained using the Adam optimiser [43]. Accurate reconstruction of the acoustic hemisphere requires the predicted field to match the simulation not only in magnitude at each observer point but also in spatial structure, directional gradients, and local intensity patterns. A conventional loss function such as RMSE or MSE evaluates the error at each observer location independently, with no awareness of neighbouring points [44]. Such a metric cannot enforce spatial coherence: two hemispheres may yield identical overall error while exhibiting very different directivity patterns. A composition loss function is therefore employed:

$$\mathcal{L}_{\mathrm{total}} = w_1\, \mathcal{L}_{\mathrm{RMSE}} + w_2\, (1 - \mathrm{SSIM}) + w_3\, \mathcal{L}_{\mathrm{grad}} + w_4\, \mathcal{L}_{\mathrm{lap}} + w_5\, \mathcal{L}_{\mathrm{wMSE}} \tag{17}$$

The five terms address distinct aspects of prediction fidelity. $\mathcal{L}_{\mathrm{RMSE}}$ minimises the overall magnitude error across all observer points on the hemisphere. The structural similarity index (SSIM) [44] evaluates whether the predicted hemisphere reproduces the spatial patterns of the simulated field, for example, the shape and extent of high-intensity lobes, rather than only matching individual point values. The gradient term $\mathcal{L}_{\mathrm{grad}}$ penalises errors in the rate of change of the OASPL between neighbouring observer points, ensuring that the predicted field captures the correct directional trends. The Laplacian term $\mathcal{L}_{\mathrm{lap}}$ penalises errors in the curvature of the field, preserving features such as the sharpness of acoustic lobes and the transitions between high- and low-intensity regions. The weighted MSE term $\mathcal{L}_{\mathrm{wMSE}}$ applies a higher penalty in regions of elevated acoustic intensity, where prediction errors have the greatest practical impact. There are no established rules for selecting the loss function weights and hyperparameters of a deep learning model, as they typically depend on the nonlinearity of the data, the size of the training set, and the dimensionality of the problem [39]. The weighting coefficients $w_1, \ldots, w_5$ and all remaining hyperparameters were determined through a trial and error process, seeking to minimise the reconstruction error across the design space; their values are reported in Section IV.B. The specific filter counts at each encoder and decoder stage are shown in Figure 4. The validation loss is monitored to prevent overfitting, and the learning rate is reduced automatically when training progress stalls.

IV. Results and Discussion

The database of $N_s = 1287$ converged rotor configurations (Sec. III.B) was partitioned into training and testing subsets to evaluate both reduced-order models on identical, independent test data. A test set of $N_{\mathrm{test}} = 129$ samples (10%) was reserved exclusively for final accuracy assessment. For the enhanced U-Net CNN, the remaining samples were split into $N_{\mathrm{train}}^{\mathrm{CNN}} = 1029$ (80%) for training and $N_{\mathrm{val}} = 129$ (10%) for validation, used to monitor generalisation performance and select model hyperparameters. For the POD-Kriging model, which does not require a separate validation set, the training and validation partitions are combined into a single training set of $N_{\mathrm{train}}^{\mathrm{PK}} = 1158$ (90%). This common test set ensures that the reconstruction accuracy of the two approaches is directly comparable. The reconstruction accuracy of both ROMs is quantified using three complementary metrics (Eq. 18): the root mean square error (RMSE, σ_f) measures the average error across all observer points, the coefficient of determination (R^2) quantifies the proportion of spatial variance captured, and the maximum absolute error (MaxAE, $\varepsilon_{\max}$) identifies the worst-case error at any single observer point.

$$\sigma_f = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\left(p_{\mathrm{sim}}^{(i)} - \hat{p}_{\mathrm{ROM}}^{(i)}\right)^2}, \qquad R^2 = 1 - \frac{\sum_{i=1}^{N}\left(p_{\mathrm{sim}}^{(i)} - \hat{p}_{\mathrm{ROM}}^{(i)}\right)^2}{\sum_{i=1}^{N}\left(p_{\mathrm{sim}}^{(i)} - \bar{p}_{\mathrm{sim}}\right)^2}, \qquad \varepsilon_{\max} = \max_{1 \le i \le N}\left|p_{\mathrm{sim}}^{(i)} - \hat{p}_{\mathrm{ROM}}^{(i)}\right|, \tag{18}$$

where $p_{\mathrm{sim}}^{(i)}$ and $\hat{p}_{\mathrm{ROM}}^{(i)}$ are the simulated and predicted OASPL at the i-th observer point, respectively, and $\bar{p}_{\mathrm{sim}}$ is the mean of the simulated field.

A. POD-Kriging Reduced-Order Model

1. Modal decomposition of the acoustic field

The POD was applied to the mean-subtracted snapshot matrix $\mathbf{S}$ of OASPL hemispheres (Eq. 1), where the mean OASPL field $\overline{\mathbf{p}}$ is computed across all $N_{\text{train}}^{\text{PK}}$ training samples and subtracted from each sample prior to decomposition. The decomposition yields an orthonormal set of spatial modes ranked by their contribution to the total acoustic variance. Figure 5 illustrates the energy distribution across the POD spectrum. The singular values (Fig. 5a) decay with increasing mode number, with the energy concentrated in the leading modes. The cumulative energy (Fig. 5b) rises steeply: Mode 1 on its own already captures 80.46% of the total variance, the first six modes account for 90.42%, and fourteen modes are required to exceed the 95% threshold (95.11%). This low-dimensional behaviour is demonstrated here for unweighted OASPL; extension to frequency-weighted metrics is identified as future work.

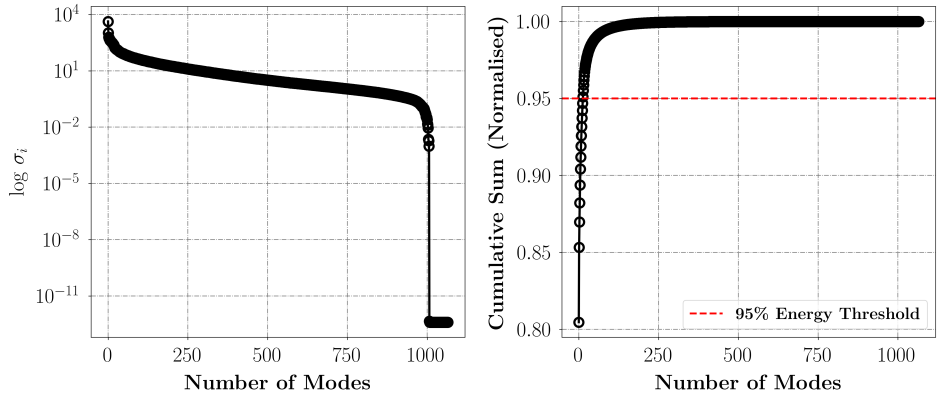


Fig. 5 Importance of POD modes: (a) decay of the POD singular values, indicating the relative strength of each mode; (b) normalised cumulative acoustic energy as a fraction of the number of retained modes.

Figure 6 presents the spatial distributions of the first twelve normalised POD modes, where the percentage in parentheses represents the cumulative acoustic energy captured when all modes up to that order are retained. The spatial patterns reveal a clear hierarchy of the dominant rotor-noise mechanisms as described in Section II. Modes 1 and 2 (80.46% and 85.33% cumulative energy, respectively) both exhibit peak amplitude near the rotor tip-path plane ($\beta \approx 0°$), consistent with the far-field directivity of thickness noise, which is governed by the monopole character of blade-volume displacement and radiates primarily in the rotor plane [3, 15]. Mode 1 shows its strongest values in a narrow band along the rotor plane, concentrated in the forward sector ($\psi \approx 180°$-$225°$), with amplitude decreasing toward lower elevations. Mode 2 exhibits a dipolar pattern, with a positive lobe at the rotor plane in the forward sector $\psi \approx 120°$-$240°$ opposing a negative region concentrated on the advancing side ($\psi \approx 30°$-$90°$). Together, these two modes account for 85.33% of the total acoustic variance, indicating that in-plane radiation is the primary source of variability across the design space. Mode 3 (1.65% individual energy, 86.98% cumulative) marks the first clear departure from the in-plane pattern of the leading two modes. Its energy is concentrated below the rotor plane, where alternating positive and negative lobes appear toward the lower hemisphere. The spatial distribution also becomes azimuthally asymmetric between the advancing ($\psi \approx 60°$) and retreating ($\psi \approx 300°$) sides. This out-of-plane, aft-dominated pattern is characteristic of loading noise, which arises from unsteady aerodynamic forces on the blade surfaces and radiates preferentially out of the rotor plane [3, 21].

Modes 4-12 capture a mixture of noise-source signatures at progressively lower energy levels. Some modes in this group retain a predominantly in-plane character reminiscent of thickness noise (e.g. Mode 6), others exhibit out-of-plane asymmetry consistent with loading noise (e.g. Mode 4), and several develop localised, high-gradient lobes concentrated in narrow azimuthal sectors (e.g. Mode 5, dominant around $\psi \approx 0$-$60°$). This interleaving of in-plane and out-of-plane features reflects the fact that POD modes are ranked by variance contribution rather than by physical mechanism, individual modes therefore combine multiple noise-source characteristics and are not exclusively attributable to a single design parameter. Modes 7-11 capture progressively finer spatial features at lower energy levels. Several of these modes exhibit narrow, high-gradient lobes at steep downward angles, consistent with the directional radiation patterns of BVI noise [5], which dominates the acoustic field in the low-speed descent flight condition simulated in this study. Although individually small, these modes are essential for resolving the sharp spatial gradients that distinguish one

design from another. Modes 13 and 14 (94.68% and 95.11% cumulative) complete the retained basis; beyond this threshold, additional modes represent higher-order corrections with diminishing energy content. Establishing a direct link between these acoustic modes and specific aerodynamic source mechanisms, for instance by applying SVD to the unsteady blade airloads, remains a direction for future investigation.

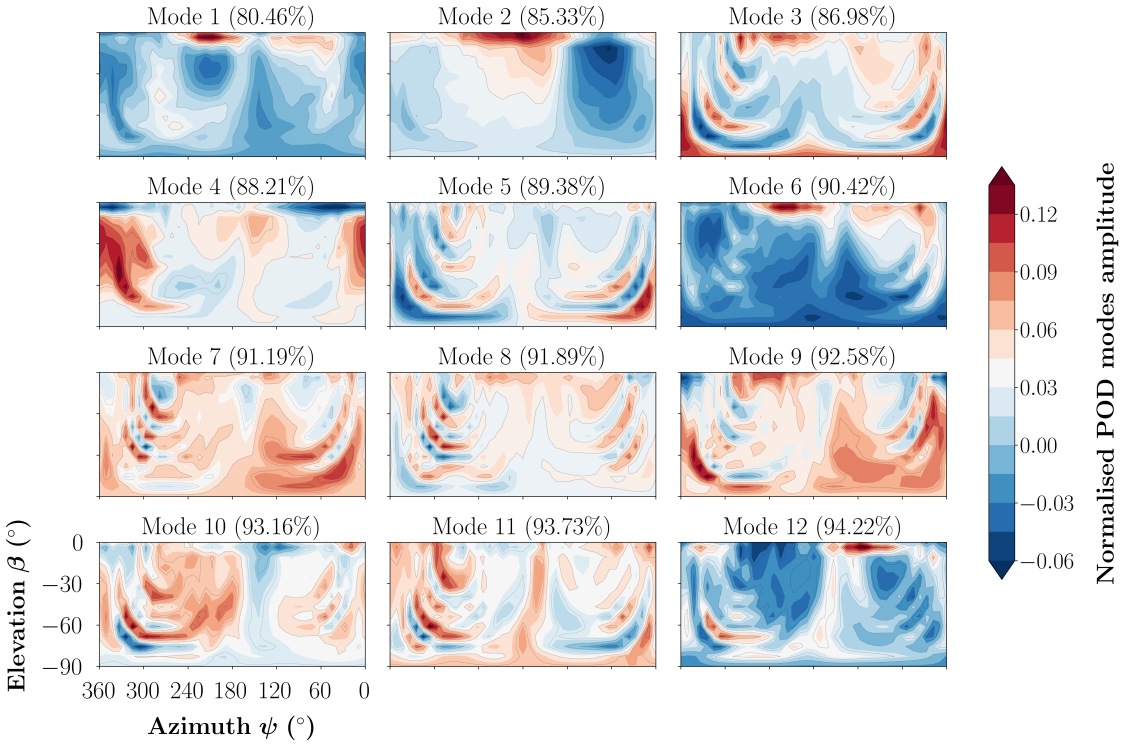


Fig. 6 Spatial distributions of the first twelve normalised POD modes of the mean-subtracted OASPL field. The percentage in parentheses indicates the cumulative acoustic energy captured up to and including that mode.

2. Kriging prediction accuracy and hemisphere reconstruction

Five-fold cross-validation was used to verify the generalisation performance of the Kriging surrogates prior to the final test-set evaluation. For each test sample, the trained Kriging models predict the modal coefficients from the design parameters, and the OASPL hemispheres are then reconstructed via the POD expansion (Eq. 14). Figure. 7 presents parity plots of the predicted versus true modal coefficients on the independent test set N_{test} for all fourteen modes. Three modes achieve excellent predictive accuracy ($R^2 > 0.90$): Mode 1 ($R^2 = 0.995$), Mode 2 ($R^2 = 0.949$), and Mode 4 ($R^2 = 0.906$), with data points clustering tightly around the 45-degree reference line. Modes 5, 6, 9, and 12 fall in the moderate range ($0.70 < R^2 \leq 0.90$), indicating moderate predictive accuracy for these intermediate modes. The remaining seven modes (3, 7, 8, 10, 11, 13, and 14) show $R^2 \leq 0.70$, reflecting the increasing difficulty of approximating the complex, nonlinear dependence of higher-order modal coefficients on the design parameters.

Despite the modest per-mode accuracy beyond Mode 4, the reconstructed OASPL hemispheres achieve a mean RMSE of 0.70 ± 0.33 dB and $R^2 = 0.986 \pm 0.015$ across the test set (Fig. 7). This apparent discrepancy is explained by the energy dominance of the leading modes: Modes 1 and 2 alone carry 85.3% of the total energy, so the well-predicted low order coefficients govern the reconstruction while the less well predicted higher-order modes contribute only minor spatial corrections. The choice of $r = 14$ retained modes reflects a deliberate trade-off between truncation fidelity and surrogate predictability. Retaining additional modes to capture 98-99% of the total energy would require approximately 31-57 modes (Fig. 5). However, the total reconstruction error comprises two competing sources: the truncation error from the omitted modes, which decreases with r, and the prediction error from Kriging inaccuracy, which grows as higher-order modes exhibit increasingly complex dependence on the design variables. Already at $r = 14$, seven of the fourteen modes have test-set R^2 values below 0.70; modes beyond the fourteenth would be predicted with even lower fidelity, effectively injecting spatially structured noise into the reconstruction. The 14-mode truncation therefore accepts

approximately 5% energy loss in exchange for reliable Kriging predictions on the retained modes. However, a systematic optimisation of the truncation level, balancing reconstruction fidelity against surrogate accuracy, is recommended for future study.

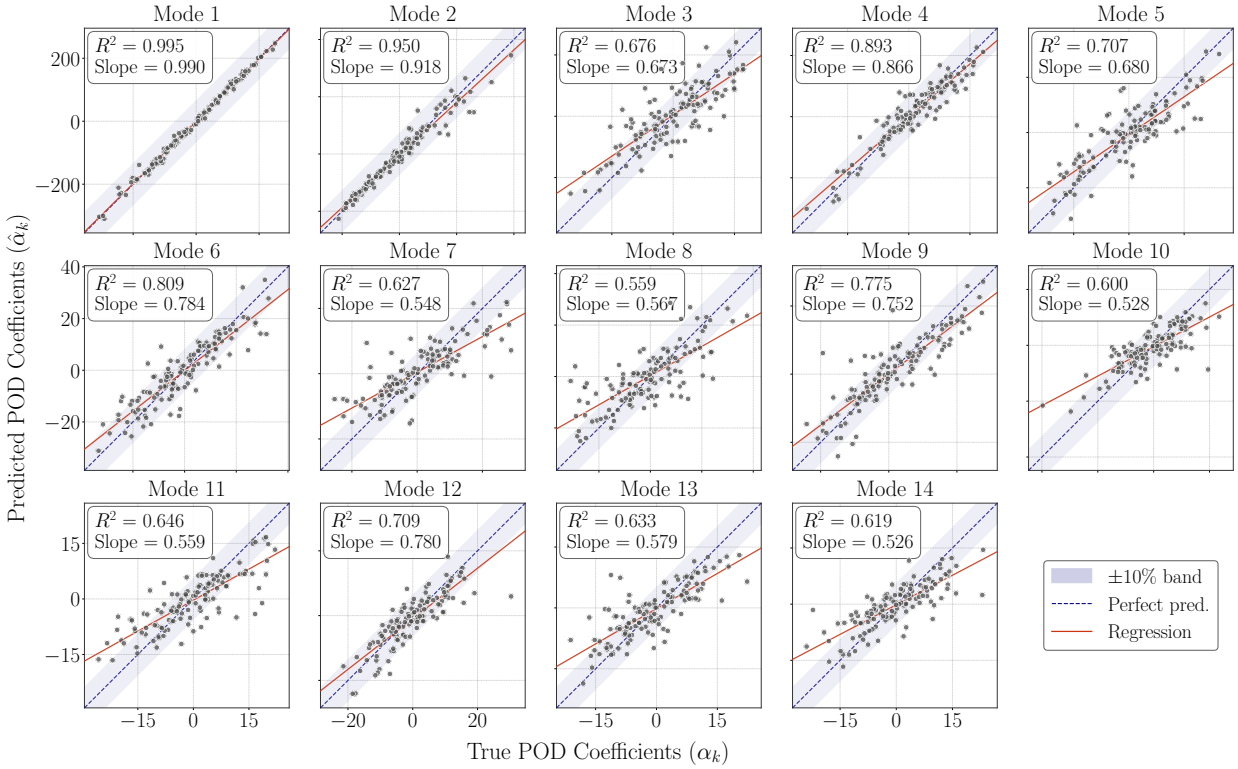


Fig. 7 Modal coefficient prediction accuracy of the POD-Kriging surrogate on independent test set ($N_{\text{test}} = 129$). Each subplot shows the predicted versus true coefficients for one POD mode.

Table 3 Design parameters for four representative test-set configurations.

Design	Ω (rad/s)	N_b (−)	c_b (m)	R (m)	λ_{taper} (−)	θ_{tw} (deg)
R1	85.3	4	0.126	2.386	0.680	5.92
R2	91.8	4	0.117	2.340	0.950	5.93
R3	108.4	5	0.106	2.004	0.928	15.76
R4	122.1	5	0.123	1.715	0.961	8.84

B. Enhanced U-Net CNN Prediction Accuracy

The model was trained on $N_{\text{train}}^{\text{CNN}}$ samples for 248 epochs using the Adam optimiser with an initial learning rate of 10^{-3} and a mini-batch size of 16. Convergence was confirmed by monitoring the validation loss on N_{val} samples, with marginal improvement observed beyond approximately 200 epochs. The composite loss weights were set to $w_1 = 0.25$ (RMSE), $w_2 = 0.25$ (SSIM), $w_3 = 0.15$ (gradient), $w_4 = 0.10$ (Laplacian), and $w_5 = 0.15$ (wMSE).

Figure 8 presents reconstructed OASPL hemispheres for four representative test set configurations (Tab. 3). The CNN successfully reproduces the dominant spatial features of the acoustic radiation: the high-intensity in-plane regions governed by thickness noise, the out-of plane asymmetry associated with loading noise, and the localised directional

lobes characteristic of BVI [3]. Across all four designs, the model achieves RMSE between 0.20 and 0.68 dB with $R^2 > 0.991$ (Tab. 4). The MaxAE ranges from 0.94 to 2.69 dB; these peaks are confined to isolated observer points in regions of steep spatial gradients, where small shifts in the predicted lobe position produce large local deviations.

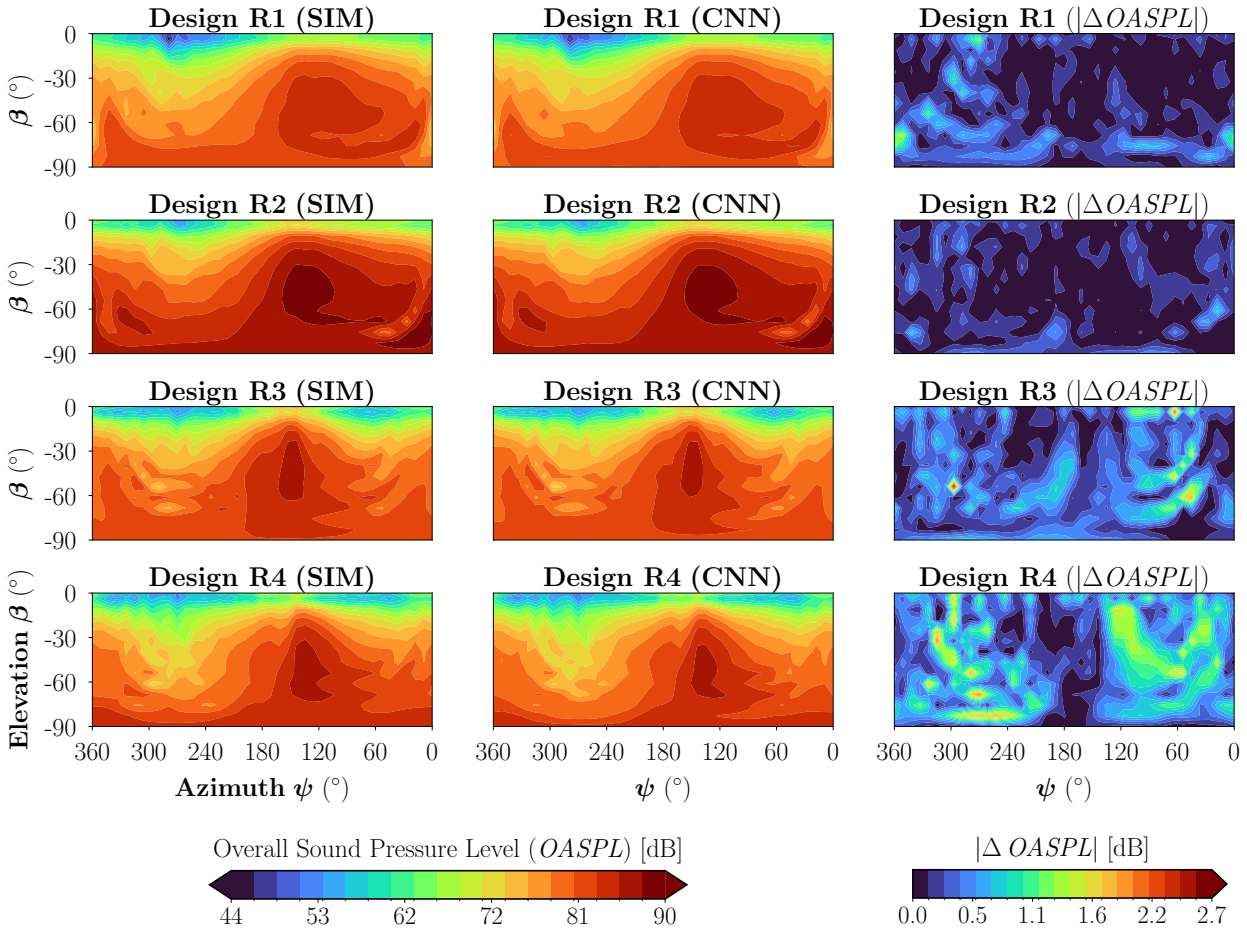


Fig. 8 OASPL hemisphere reconstruction by the enhanced U-Net CNN for designs R1-R4 (Tab. 3).

C. Comparative Assessment of Reduced-Order Models

Table 4 compares the reconstruction accuracy of both models on the four representative test set configurations (Tab. 3). The enhanced U-Net CNN consistently outperforms the POD-Kriging model across all three metrics. The CNN reduces both RMSE and MaxAE by a factor of approximately 2-4 relative to POD-Kriging. The reconstruction accuracy generally decreases for both models as Ω increases and R decreases, corresponding to higher M_{tip} and more complex spatial gradients in the acoustic field. The largest disparity between the two models occurs for design R3, where the POD-Kriging MaxAE is 7.68 dB compared with 2.69 dB for the CNN (Tab. 4).

Figure 9 compares the simulated and reconstructed hemispheres for all four designs. The columns show the numerical simulation (Sec. III.A), the POD-Kriging and CNN predictions. For designs R1 and R2 both models reproduce the simulation closely. For designs R3 and R4, where the acoustic field exhibits more complex spatial structure in the lower hemisphere, the differences between the two models become apparent: the POD-Kriging reconstruction tends to smooth or distort localised features below the rotor plane, whereas the CNN preserves the spatial detail of the simulated field more accurately. Figure 10 presents predicted versus simulated OASPL values for all observer points across the the R1-R4 configurations. The CNN predictions cluster more tightly around the identity line (i.e., 45° line), whereas the POD-Kriging predictions show wider dispersion, particularly in the high OASPL region where the sharpest spatial gradients occur.

Table 4 Reconstruction accuracy of ROMs for four representative test-set configurations (Tab. 3).

	POD-Kriging			Enhanced U-Net CNN		
Design	RMSE (dB)	MaxAE (dB)	R^2 (−)	RMSE (dB)	MaxAE (dB)	R^2 (−)
R1	0.88	5.25	0.984	0.31	1.61	0.998
R2	0.80	4.10	0.986	0.20	0.94	0.999
R3	1.05	7.68	0.974	0.50	2.69	0.994
R4	1.34	4.52	0.965	0.68	2.10	0.991

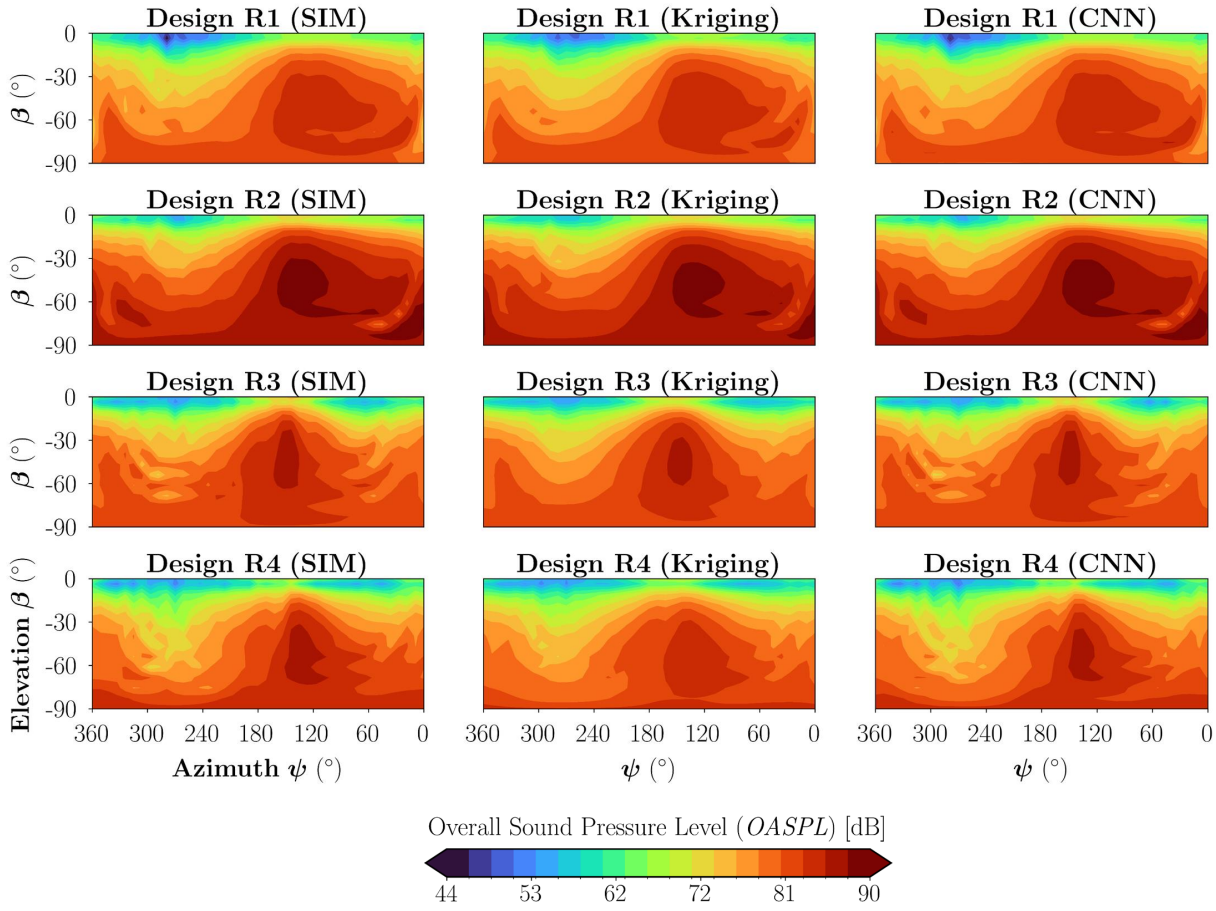


Fig. 9 OASPL hemisphere comparisons between POD-Kriging and CNN for designs R1-R4 (Tab. 3).

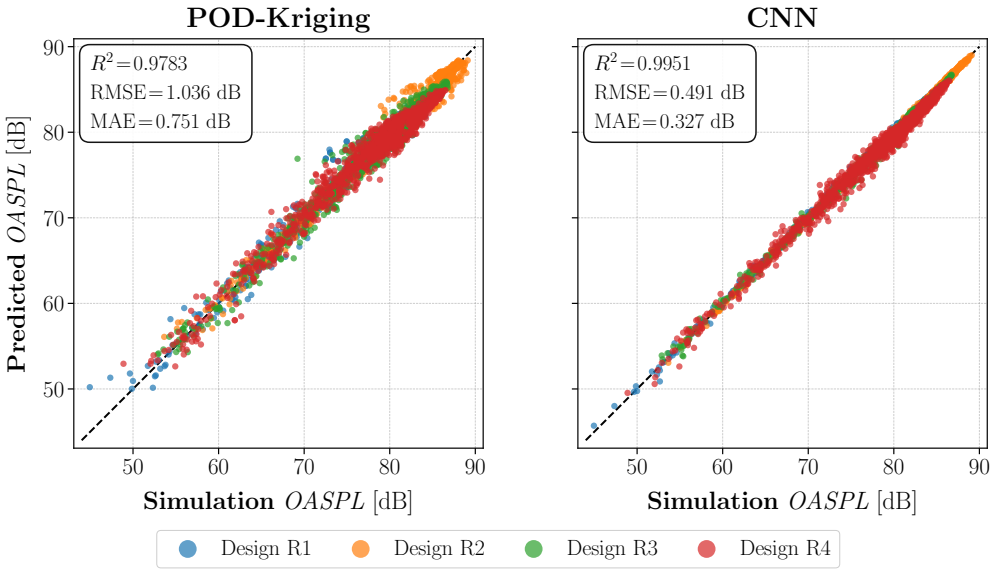
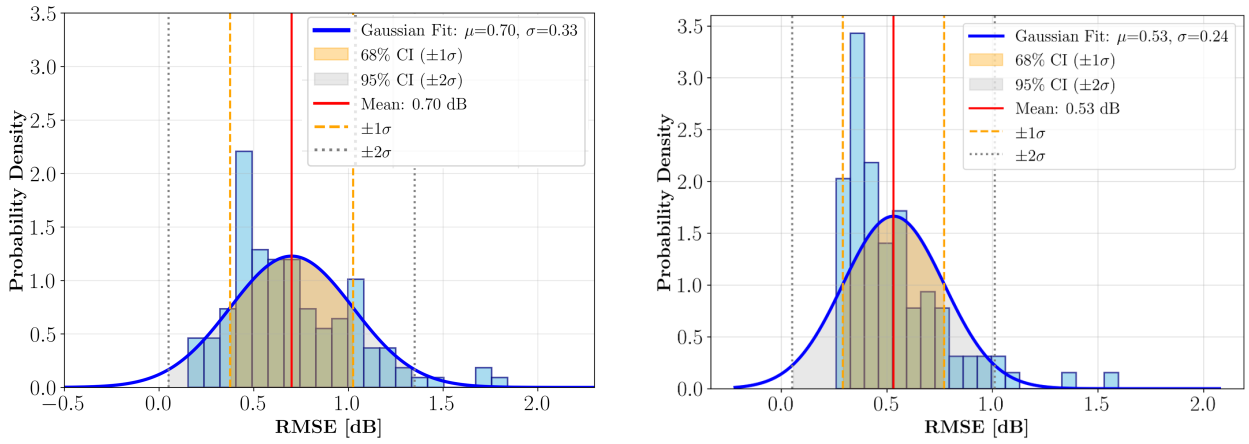


Fig. 10 Comparison of predicted versus simulated OASPL of both ROMs for designs R1-R4 (Tab. 3).

Figure 11 compares the RMSE distributions across all N_{test} test samples. The CNN achieves a lower mean (μ = 0.53 dB versus μ = 0.70 dB) and tighter spread (σ = 0.24 dB versus σ = 0.33 dB), indicating higher accuracy and greater consistency across the design space.



(a) POD-Kriging **(b) U-CNN**

Fig. 11 RMSE distribution across all N_{test} test samples with Gaussian fit for (a) POD-Kriging and (b) enhanced U-Net CNN.

Two factors may contribute to the improved accuracy of the CNN. First, the CNN operates directly on the two-dimensional hemisphere grid, preserving spatial correlations through convolutional filters and the composite loss function (Eq. 17), which explicitly penalises errors in spatial structure, gradients, and curvature. The POD-Kriging model, by contrast, predicts each modal coefficient independently, and the reconstruction accuracy is constrained by both the 14-mode truncation and the Kriging fidelity on the higher-order modes (Sec. IV.A.2). Second, the CNN learns a direct mapping from design parameters to the full hemisphere in a single forward pass, whereas POD-Kriging requires r separate surrogate models whose individual prediction errors combine during reconstruction.

The POD-Kriging model offers two advantages over the CNN. First, the POD decomposition provides a ranked set of orthonormal modal distributions with quantified energy contributions (Sec. IV.A.1), enabling the designer to identify which spatial patterns dominate the acoustic field and to select the number of retained modes for a desired fidelity level.

Second, the Kriging surrogates require no iterative hyperparameter tuning, validation set, or composite loss function design, all model parameters are estimated analytically by maximum likelihood estimation. Both models evaluates a single hemisphere in under 1 s, compared with approximately 2 hours per configuration for the full simulation chain, confirming that both approaches provide practical reduced-order alternatives for preliminary acoustic assessment.

V. Conclusions

Two reduced-order models (ROMs) have been developed for the direct prediction of spatially resolved OASPL acoustic hemispheres from rotor design parameters: a POD-Kriging model combining proper orthogonal decomposition with Kriging regression, and an enhanced U-Net convolutional neural network. A database of $N_s = 1287$ converged rotor configurations was generated using a validated free-wake aeroelastic and acoustic simulation chain, across a six-dimensional design space around the HART II baseline rotor in a descent flight condition. The POD decomposition reveals that the acoustic field is effectively low-dimensional, with 14 modes capturing 95.11% of the total variance. The POD-Kriging model achieves a mean RMSE = 0.70 dB and $R^2 = 0.986$ across the $N_{\text{test}} = 129$ independent test samples. The enhanced U-Net CNN achieves a mean RMSE = 0.53 dB across the same test set. On the four representative designs (Tab. 4), the CNN outperforms POD-Kriging by a factor of 2-4 in both RMSE and MaxAE. Both models achieve mean RMSE < 1.0 dB with mean $R^2 > 0.96$ across the test set.

The novelty of this work is the compilation of a large-scale dataset of aeroacoustic noise hemispheres across a wide range of rotor designs and operating conditions, and the synthesis of two ROMs capable of predicting the acoustic behaviour of the investigated rotor configurations in a reduced manner. This is an advance beyond the previous state of the art, as existing models do not take rotor design parameters into account. These parameters have been included in this work, enabling acoustic assessment across the rotor design space. The contribution is that the work proves that POD-Kriging reconstruction accuracy is constrained by the predictability of higher-order modal coefficients, leading to degraded performance in regions of steep directional gradients, whereas the CNN bypasses this limitation by learning the full spatial field directly at the cost of requiring careful hyperparameter and loss function design. These findings provide practical modelling guidelines for ROM selection in rotorcraft acoustic applications. The impact of this work is the developed approach, which enables a reduction in the computational time required for noise hemisphere estimation by approximately four orders of magnitude relative to the underlying solver chain. This allows the introduction of aeroacoustic assessment with sufficient predictive pedigree into the preliminary design stage of rotorcraft, where aeroacoustic performance is typically absent from the design trade space.

There are several promising avenues for future work. For the POD-Kriging model, increasing the number of retained modes and selecting the correlation function on a per-mode basis may improve reconstruction fidelity; both directions involve a trade-off between accuracy and computational overhead. For the CNN, systematic hyperparameter optimisation and a rigorous ablation study are recommended. For both models, adaptive enrichment of the training database in regions of the design space with the largest reconstruction errors is expected to improve accuracy. Longer-term objectives include extension to multiple flight conditions, interactional noise between main and tail rotors, and establishing a direct link between the POD modes and rotor design parameters.

Acknowledgments

The authors would like to express their gratitude to the Engineering and Physical Sciences Research Council (EPSRC) Grant EP/X52475X/1 and the Defence Science and Technology Laboratory (Dstl) for funding this project and granting permission to publish this research. The authors also thank Richard Markiewicz and Neil Taylor of the Rotorcraft Analysis and Engineering Team at Dstl for their guidance and support.

References

[1] UK Ministry of Defence, "Defence Aviation Net Zero Strategy," , 2023.

[2] European Commission, "Flightpath 2050: Europe's Vision for Aviation," Tech. rep., Publications Office of the European Union, Luxembourg, 2011. https://doi.org/10.3233/978-1-61499-063-5-63.

[3] Brentner, K. S., and Farassat, F., "Modeling Aerodynamically Generated Sound of Helicopter Rotors," *Progress in Aerospace Sciences*, Vol. 39, No. 2–3, 2003, pp. 83–120. https://doi.org/10.1016/S0376-0421(02)00068-4.

[4] van der Wall, B. G., Lim, J. W., Baeder, J. D., and Boyd, D. D., "The HART II International Workshop: An Assessment

of the State-of-the-Art in Comprehensive Code Prediction," *CEAS Aeronautical Journal*, Vol. 4, No. 3, 2013, pp. 223–252. https://doi.org/10.1007/s13272-013-0077-9.

[5] Leishman, J. G., *Principles of Helicopter Aerodynamics*, 2nd ed., Cambridge University Press, 2016.

[6] Thurman, C. S., and Zawodny, N. S., "Aeroacoustic Characterization of Optimum Hovering Rotors Using Artificial Neural Networks," Tech. Rep. 20210013432, NASA Langley Research Center, 2021. URL https://ntrs.nasa.gov/citations/20210013432.

[7] Li, S., and Lee, S., "A Machine Learning-Based Fast Prediction of Rotorcraft Broadband Noise," *AIAA AVIATION 2020 Forum*, 2020. https://doi.org/10.2514/6.2020-2588, aIAA Paper 2020-2588.

[8] Mankbadi, R. R., Afari, S., and Golubev, V. V., "Development of New Empirical Rotor Broadband Noise Prediction Models for Urban Air Mobility Applications," *AIAA SciTech 2024 Forum*, 2024. https://doi.org/10.2514/6.2024-2473, aIAA Paper 2024-2473.

[9] Thurman, C. S., and Zawodny, N. S., "An Artificial Neural Network Approach to Predict Rotor-Airframe Acoustic Waveforms," *AIAA AVIATION 2023 Forum*, 2023. https://doi.org/10.2514/6.2023-4182, aIAA Paper 2023-4182.

[10] Wiedemann, A. D., Fuller, C., and Pascioni, K. A., "Constructing a Physics-Guided Machine Learning Neural Network to Predict Tonal Noise Emitted by a Propeller," *INTER-NOISE and NOISE-CON Congress and Conference Proceedings*, Vol. 264, 2022, pp. 151–162. https://doi.org/10.3397/NC-2022-709.

[11] Zhou, B. Y., Hanson, L., Pullin, S. F., Zang, B., Hauth, J., and Huan, X., "Data-Driven Propeller Noise Prediction Approach with Transfer and Active Learning," *AIAA Journal*, Vol. 63, No. 12, 2025. https://doi.org/10.2514/1.J065406.

[12] Gennaretti, M., Serafini, J., Bernardini, G., Castorrini, A., De Matteis, G., and Avanzini, G., "Numerical Characterization of Helicopter Noise Hemispheres," *Aerospace Science and Technology*, Vol. 52, 2016, pp. 18–28. https://doi.org/10.1016/j.ast.2016.02.013.

[13] Vouros, S., Goulos, I., and Pachidis, V., "Integrated Methodology for the Prediction of Helicopter Rotor Noise at Mission Level," *Aerospace Science and Technology*, Vol. 89, 2019, pp. 136–149. https://doi.org/10.1016/j.ast.2019.03.061.

[14] Kwon, S., Hong, Y., and Yee, K., "Prediction of Ground Noise Footprint Using Reduced Order Model," *AIAA AVIATION 2023 Forum*, 2023. https://doi.org/10.2514/6.2023-3654, aIAA Paper 2023-3654.

[15] Farassat, F., "Derivation of Formulations 1 and 1A of Farassat," Tech. Rep. TM-214853, NASA, 2007.

[16] Goulos, I., Giannakakis, P., Pachidis, V., and Pilidis, P., "Mission performance simulation of integrated helicopter–engine systems using an aeroelastic rotor model," *Journal of Engineering for Gas Turbines and Power*, Vol. 135, No. 9, 2013, p. 091201. https://doi.org/10.1115/1.4024869.

[17] Goulos, I., Pachidis, V., and Pilidis, P., "Helicopter rotor blade flexibility simulation for aeroelasticity and flight dynamics applications," *Journal of the American Helicopter Society*, Vol. 59, No. 4, 2014, pp. 1–18. https://doi.org/10.4050/JAHS.59.042006.

[18] Goulos, I., Pachidis, V., and Pilidis, P., "Lagrangian formulation for the rapid estimation of helicopter rotor blade vibration characteristics," *The Aeronautical Journal*, Vol. 118, No. 1206, 2014, pp. 861–901. https://doi.org/10.1017/S000192400000960X.

[19] Goulos, I., "An improved analytical approach for modeling the effect of rotor wake curvature using finite-state induced flow models," *Journal of the American Helicopter Society*, Vol. 61, No. 3, 2016, pp. 1–16. https://doi.org/10.4050/JAHS.61.032010.

[20] MacMillan, W. L., "Development of a Modular-Type Computer Program for the Calculation of Gas Turbine Off-Design Performance," , 1974.

[21] Vouros, S., "Aeroacoustic Simulation of Rotorcraft Propulsion Systems," Ph.d. thesis, Cranfield University, Cranfield, Bedfordshire, UK, 2019.

[22] Castillo Pardo, A., "Aeroelastic Simulation of Rotorcraft Propulsion Systems," Ph.d. thesis, Cranfield University, Cranfield, Bedfordshire, UK, 2017.

[23] Ffowcs Williams, J. E., and Hawkings, D. L., "Sound Generation by Turbulence and Surfaces in Arbitrary Motion," *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, Vol. 264, No. 1151, 1969, pp. 321–342. https://doi.org/10.1098/rsta.1969.0031.

[24] Pegg, R., "A Summary and Evaluation of Semi-Empirical Methods for the Prediction of Helicopter Rotor Noise," Technical Memorandum NASA TM-80200, NASA, 1979.

[25] Snider, R., Samuels, T., Goldman, B., and Brentner, K. S., "Full-Scale Rotorcraft Broadband Noise Prediction and Its Relevance to Civil Noise Certification Criteria," *Proceedings of the 69th Annual Forum of the American Helicopter Society*, American Helicopter Society, Phoenix, AZ, 2013.

[26] van der Wall, B. G., Burley, C. L., Yu, Y., Richard, H., Pengel, K., and Beaumier, P., "The HART II Test: Measurement of Helicopter Rotor Wakes," *Aerospace Science and Technology*, Vol. 8, No. 4, 2004, pp. 273–284. https://doi.org/10.1016/j.ast.2004.01.001.

[27] Lim, J. W., van der Wall, B. G., Baeder, J. D., Boyd, D. D., and Burley, C. L., "HART II: Prediction of Blade-Vortex Interaction Loading," *Proceedings of the 29th European Rotorcraft Forum (ERF)*, Friedrichshafen, Germany, 2003, pp. 16–18.

[28] Schneider, O., "Analysis of SPR Measurements from HART II," *Aerospace Science and Technology*, Vol. 9, No. 5, 2005, pp. 409–420. https://doi.org/10.1016/j.ast.2005.01.013.

[29] van der Wall, B. G., "A Comprehensive Rotary-Wing Data Base for Code Validation: The HART II International Workshop," *The Aeronautical Journal*, Vol. 115, No. 1164, 2011.

[30] Olsson, A., Sandberg, G., and Dahlblom, O., "On Latin Hypercube Sampling for Structural Reliability Analysis," *Structural Safety*, Vol. 25, No. 1, 2003, pp. 47–68. https://doi.org/10.1016/S0167-4730(02)00039-5.

[31] Taira, K., Brunton, S. L., Dawson, S. T. M., Rowley, C. W., Colonius, T., McKeon, B. J., Schmidt, O. T., Gordeyev, S., Theofilis, V., and Ukeiley, L. S., "Modal Analysis of Fluid Flows: An Overview," *AIAA Journal*, Vol. 55, No. 12, 2017, pp. 4013–4041. https://doi.org/10.2514/1.J056060.

[32] Brunton, S. L., and Kutz, J. N., *Data-Driven Fluid Mechanics: Combining First Principles and Machine Learning*, Cambridge University Press, 2022. https://doi.org/10.1017/9781009089517.

[33] Simpson, T. W., Mauery, T. M., Korte, J. J., and Mistree, F., "Kriging models for global approximation in simulation-based multidisciplinary design optimization," *AIAA Journal*, Vol. 39, No. 12, 2001, pp. 2233–2241. https://doi.org/10.2514/2.1234.

[34] Bouhlel, M. A., Hwang, J. T., Bartoli, N., Lafage, R., Morlier, J., and Martins, J. R. R. A., "A Python Surrogate Modeling Framework with Derivatives," *Advances in Engineering Software*, Vol. 135, 2019, p. 102662. https://doi.org/10.1016/j.advengsoft.2019.03.005.

[35] Ronneberger, O., Fischer, P., and Brox, T., "U-Net: Convolutional Networks for Biomedical Image Segmentation," *arXiv preprint arXiv:1505.04597*, 2015. https://doi.org/10.48550/arXiv.1505.04597.

[36] Zhang, Y., Tian, Y., Kong, Y., Zhong, B., and Fu, Y., "Residual Dense Network for Image Super-Resolution," *arXiv preprint arXiv:1802.08797*, 2018. https://doi.org/10.48550/arXiv.1802.08797.

[37] Oktay, O., Schlemper, J., Folgoc, L. L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N. Y., Kainz, B., Glocker, B., and Rueckert, D., "Attention U-Net: Learning Where to Look for the Pancreas," *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2018. https://doi.org/10.48550/arXiv.1804.03999.

[38] Woo, S., Park, J., Lee, J.-Y., and Kweon, I. S., "CBAM: Convolutional Block Attention Module," *arXiv preprint arXiv:1807.06521*, 2018. https://doi.org/10.48550/arXiv.1807.06521.

[39] Goodfellow, I., Bengio, Y., and Courville, A., *Deep Learning*, MIT Press, 2016. URL https://www.deeplearningbook.org.

[40] LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P., "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, Vol. 86, No. 11, 1998, pp. 2278–2324. https://doi.org/10.1109/5.726791.

[41] Tejero, F., Sureshbabu, S., Boscagli, L., and MacManus, D., "Point-enhanced convolutional neural network: A novel deep learning method for transonic wall-bounded flows," *Aerospace Science and Technology*, Vol. 155, 2024, p. 109689. https://doi.org/10.1016/j.ast.2024.109689.

[42] Simonyan, K., and Zisserman, A., "Very Deep Convolutional Networks for Large-Scale Image Recognition," *arXiv preprint arXiv:1409.1556*, 2015. https://doi.org/10.48550/arXiv.1409.1556.

[43] Kingma, D. P., and Ba, J., "Adam: A Method for Stochastic Optimization," *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015. https://doi.org/10.48550/arXiv.1412.6980.

[44] Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P., "Image quality assessment: from error visibility to structural similarity," *IEEE Transactions on Image Processing*, Vol. 13, No. 4, 2004, pp. 600–612. https://doi.org/10.1109/TIP.2003.819861.



https://dspace.org/

Machine learning approaches for preliminary design of rotorcraft aeroacoustics

Teerawathananon, Peekarn; Goulos, Ioannis; MacManus, David G.

2026-05-26
American Institute of Aeronautics and Astronautics (AIAA)
https://doi.org/10.2514/6.2026-3451
Teerawathananon P, Goulos I, MacManus D. (2026) Machine learning approaches for preliminary design of rotorcraft aeroacoustics. In: The Proceedings of the 32nd AIAA/CEAS Aeroacoustics Conference (2026), 26-29 May 2026, Brussels, Belgium, Article number 2026-3451

Downloaded from a DSpace repository. https://dspace.org/