



# A rotorcraft in-flight ice detection framework using computational aeroacoustics and Bayesian neural networks

Myles Morelli<sup>1</sup> · Jeremiah Hauth<sup>2</sup> · Alberto Guardone<sup>1</sup> · Xun Huan<sup>2</sup> · Beckett Y. Zhou<sup>3</sup> 

Received: 6 September 2022 / Revised: 7 May 2023 / Accepted: 22 May 2023 / Published online: 21 August 2023  
© The Author(s) 2023

## Abstract

This work develops a novel ice detection framework specifically suitable for rotorcraft using computational aeroacoustics and Bayesian neural networks. In an offline phase of the work, the acoustic signature of glaze and rime ice shapes on an oscillating wing are computed. In addition, the aerodynamic performance indicators corresponding to the ice shapes are also monitored. These performance indicators include the lift, drag, and moment coefficients. A Bayesian neural network is subsequently trained using projected Stein variational gradient descent to create a mapping from the acoustic signature generated by the iced wings to predict their performance indicators along with quantified uncertainty that is highly important for time- and safety-critical decision-making scenarios. While the training is carried out fully offline, usage of the Bayesian neural network to make predictions can be conducted rapidly online allowing for an ice detection system that can be used in real time and in-flight.

**Keywords** Ice detection system · Aircraft icing · Aeroacoustics · Neural networks · Bayesian

## 1 Introduction

Rotorcraft possess the capability to vertically take-off and land, allowing them to operate in challenging flight conditions where conventional fixed-wing aircraft cannot. Their operational environment lends itself to low-level altitudes,

and many missions are often toward the boundary of the flight envelope. Throughout the course of aviation, the requirement for flight in poor weather conditions and at night has always been a driver for innovation and rotorcraft are no exception. In the past, rotorcraft were fair weather vehicles with marginal performance; however, now they regularly operate in conditions from hot and dry to cold, wet, and windy (Padfield 2008). In particular, cold climates can potentially leave aircraft susceptible to dangerous icing conditions. In-flight icing encounters can jeopardize the performance and handling qualities of aircraft and hence pose a serious threat to flight safety (Bragg et al. 1986; Gent et al. 2000). This threat to flight safety arises as ice accretion can rapidly alter aerodynamic lifting surfaces such as wings and rotors during flight which are highly sensitive to geometric modifications. Naturally, establishing a comprehensive understanding of the fundamental physics of aircraft icing is paramount to ensure the reduction of ice-related accidents and the progression of safety-critical technology.

In 2002, Bragg et al. (2022) critically assessed aircraft safety during icing conditions and introduced smart icing systems to reduce aircraft accident rates. Their work recognized that warning the pilot of ice conditions through the use of ice detection systems could significantly increase the safety of aircraft to either optimize the use of ice protection

---

Responsible Editor: Chao Hu

---

Jeremiah Hauth, Alberto Guardone, Xun Huan, and Beckett Y. Zhou have contributed equally to this work.

---

Topical Collection: Advanced Optimization Enabling Digital Twin Technology  
Guest Editors: C Hu, VA González, T Kim, O San, and Z Hu

---

✉ Beckett Y. Zhou  
beckett.zhou@bristol.ac.uk

- <sup>1</sup> Department of Aerospace Sciences and Technologies, Politecnico di Milano, Building B12 Campus Bovisa, Via La Masa, 34, 20156 Milan, Italy
- <sup>2</sup> Department of Mechanical Engineering, University of Michigan, 1231 Beal Ave, Ann Arbor, MI 48109-2133, USA
- <sup>3</sup> Department of Aerospace Engineering, University of Bristol, Queens Building, University Walk, Bristol BS8 1TR, UK

systems or to avoid icing conditions altogether. Over the past two decades, this has led to sustained technological advancements in ice detection systems.

Methods for ice detection on aircraft can be characterized as either direct or indirect. The former involves using sensors to detect ice accretion on the airframe or atmospheric conditions symptomatic of icing environments, while the latter relies on monitoring for a change in response to the aircraft state due to icing. Direct methods have received the most scientific interest where many novel sensing techniques have been developed. One prominent direct strategy entails measuring the reaction of a signal or electric field to the presence of ice, such as through a surface-generated wave, infrared thermography, optical reflectometry, or ultrasonics wave techniques. For example, Varadan et al. (1997) proposed producing surface-generated Love waves from piezoelectric devices and propagating them through a surface. As ice accretes, the amplitude and velocity of these waves would be altered. Hongerholt et al. (2002) propagated and reflected ultrasonic-guided waves on the surface of a wing and measured the attenuated amplitude of the reflected waves in different mediums for detecting the existence of ice. Zhuge et al. (2012) used infrared thermography where emitted electromagnetic signals reflected by the surface are analyzed for temperature variations. Bassey and Simpson (2007) employed optical time domain reflectometry to measure the dielectric permittivity and volumetric water content on the aircraft surface. Schlegl et al. (2015) further used capacitive techniques to monitor the presence of dielectric substances, leveraging the different electrical properties of air, water, and ice. Roy et al. (1998) generated an excitation signal and measured the impedance change in the electrodes when ice is formed.

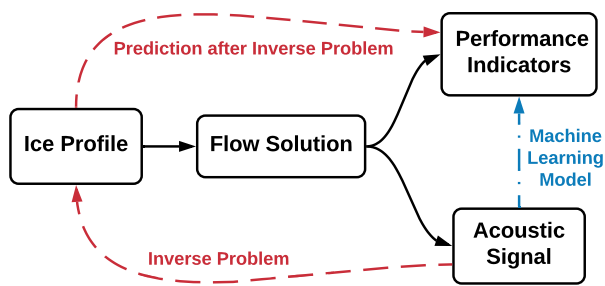
Overall, direct ice detection techniques allow for precise information regarding the ice location, thickness, and location. However, they are largely commercially unproven due to preliminary testing not exceeding the laboratory environment or are unpractical for moving components, making their implementation on rotorcraft challenging.

Indirect ice detection approaches are based on comparing the aircraft state to previously established clean states and using the difference between the iced and clean aircraft states as an indicator for icing. Melody et al. (2000) utilized aircraft dynamics as an indicator for icing, which required algorithms to identify changes in aircraft performance, stability, and control parameters over time based on measurements of the aircraft state variables and control input. Deiler and Fezans (2020) introduced a performance-based ice detection methodology based on the change in the steady flight state, which can be used prior to reaching dynamic stall conditions and during steady flight conditions contrary to the aircraft dynamics approach. However, a major limitation of using either the dynamics or performance of aircraft

for ice detection is that the type, location, and thickness of ice cannot be determined. Additionally, the detection of ice prior to dynamic stall is fundamental to aircraft safety.

Another major type of indirect ice detection involves using acoustic noise. The acoustics-based techniques are particularly relevant for rotorcraft, which bear a characteristic and largely undesirable feature that is their rotor noise. Rotor noise sources are complex and are a combination of loading noise, thickness noise, shock noise, blade–vortex interaction noise, and broadband noise. However, Cheng et al. (2018) recognized the strong correlation between the ice-induced surface roughness and the broadband noise level suggesting rotor acoustics could be a viable technology for rotorcraft ice detection. Chen et al. (2019) further investigated using numerical techniques and an artificial iced notch and concurred that the acoustics could be potentially used to indicate the formation of ice on a rotor. The experimental work from Cheng et al. (2018) used sandpaper to represent ice roughness, while the numerical work from Chen et al. (2019) used an artificial ice notch based on the experimental iced surface roughness results of a straight wing from the University of Illinois, presented by Broeren and Bragg (2005). The first characterization between glaze and rime iced noise signals was established by Morelli et al. (2020). Their work utilized numerical techniques to simulate the accretion of ice on rotors and subsequently the rotor noise signature. The results highlighted that glaze ice accretion produces noise at a significantly higher frequency caused by ice-induced flow separation. In the main, ice detection via means of aircraft acoustics is still in the early stages of research.

This paper seeks to present a proof-of-concept acoustics-based ice detection system for rotorcraft through building a machine learning (ML) model. While several studies mentioned above have investigated using noise as a method for ice detection, this work is the first to develop a system that predicts the aerodynamic performance (e.g., lift and drag coefficients) directly from observed acoustics. In other words, we seek to not only state whether icing is present or not, but also predict the severity and consequence of the icing condition through the aerodynamic performance metrics directly useful for aviation and piloting. To achieve this, we first utilize numerical techniques to simulate a large dataset of ice shapes. Next, high-fidelity scale-resolving methods are used to solve for the flow field over the iced rotor blades and extract their corresponding aerodynamic performance indicators. The surface pressure fluctuations are then taken from the aerodynamic data and used by an aeroacoustic solver to predict the far-field noise levels. Finally, we train a ML model to directly map the acoustic noise signal produced from the ice structures to the aerodynamic performance indicators of the aircraft; this “shortcut” can allow



**Fig. 1** Workflow routes illustrating the physics causation (black solid), a traditional inverse problem route that needs to first identify the ice profile from acoustic signal and then solve for performance indicators (dashed red), and proposed ML model “shortcut” mapping (dotted-dash blue). (Color figure online)

predictions to be made at real-time speeds. An overview of our overall framework is shown in Fig. 1.

We choose neural network (NN) as our ML model for its expressiveness in capturing complex non-linear mappings and in particular, Bayesian neural network (BNN) that also provides uncertainty quantification (UQ) in the ML model and its predictions. Other ML model types may also be adopted: for example, Gaussian process is another popular choice with built-in UQ, but it lacks scalability (covariance matrix inversions become expensive when more training data become available) and flexibility (typically limited for predicting low-dimensional quantities). In contrast, NNs are known for their expressiveness and ability to handle large datasets. Furthermore, the BNN framework can flexibly couple with advanced NN architectures such as convolutional neural networks (CNN) (Lecun et al. 1998) and long short-term memory units (LSTM) (Hochreiter and Schmidhuber 1997) that are specialized in high-dimensional data and time series signals, respectively. Indeed, NNs have already been used for in-flight ice detection and identification (Caliskan and Hajiyev 2013) albeit not with rotorcraft acoustics. We further highlight the critical need to provide UQ for the ML models used in such mission- and safety-critical settings—we need to know the *quality* of a ML prediction before using it for decision-making. However, performing UQ for a NN is a challenging task due to its high number of model parameters. A further contribution of our work is then to enable UQ computation in BNNs leveraging recently developed algorithm of projected Stein variational gradient descent (pSVGD) (Chen and Ghattas 2020).

Several limitations of this complex problem do however need to be acknowledged. While this methodology is able to predict the tonal and broadband noise components of iced airfoils, in reality the noise of an operating helicopter is highly complex with many different components including engine noise and tail rotor noise which would likely pollute the small differences in noise observed by an iced rotor.

Nevertheless, this is the first step as to understanding if an acoustic ice detection system using NN is feasible. To that end, if successful the installation on helicopter in operation would be possible due to the system being non-intrusive.

In summary, we develop a framework for rotorcraft icing detection highlighted by the following key novelty and contributions:

- Uses in-flight *acoustic* measurements to enable prediction of aerodynamic performance for rotorcraft, who cannot use visual/image observations due to the high-blade RPM during operation;
- Enables *real-time* performance predictions with *quantified uncertainty* through advanced ML models of BNNs, using a recently developed pSVGD algorithm;
- Uses simulation data that are generated from state-of-the-art *physics-based models* rather than empirical models; and
- Requires a truly multi-disciplinary approach to combine ice formation and growth, turbulent flow aerodynamics, aeroacoustics, and ML with UQ.

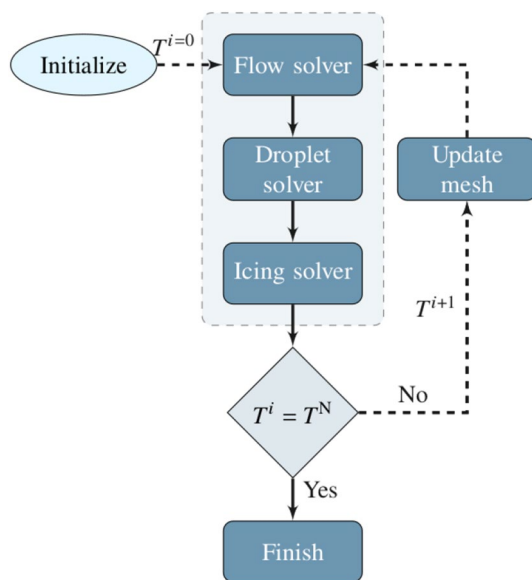
The structure of the remaining paper is as follows: Sect. 2 introduces the methodology of the acoustic ice detection and the ML using BNNs; Sect. 3 discusses the results of the ice prediction, the turbulent flow simulations and acoustic prediction, as well as the training of the BNN and the design of the ice detection system; and Sect. 4 highlights the concluding remarks of the work.

## 2 Methodology

The development of a technique for real-time, in-flight aeroacoustic ice detection system requires a highly inter-disciplinary approach. The methodology used for the prediction of in-flight ice accretion, the simulation of the turbulent iced flow field and noise prediction, and finally the ML are hereinafter introduced.

### 2.1 Icing simulation

The PoliMice ice prediction toolkit is utilized to simulate a large sample of ice shapes required for the ML model training. The PoliMice ice prediction framework is described in detail in Gori et al. (2015a) and Morelli and Guardone (2021) and will subsequently be summarized. A conventional icing simulation structure routinely involves a three-stage process which iteratively updates to account for unsteady ice accretion. This simulation process is commonly known as multi-step ice accretion and is adopted by PoliMice in this work. A schematic of multi-step ice accretion is shown in Fig. 2.



**Fig. 2** Multi-step ice accretion simulation

The first stage of the icing framework is to simulate the aerodynamic flow field around the region of interest. In this work, the SU2 solver is used to compute the aerodynamic flow field over the wing, with further details of the SU2 solver discussed in Sect. 2.2.

The second stage is to simulate the trajectories of supercooled water droplets within the flow domain and determine the collection efficiency and impingement locations. An in-house particle tracking code based on a Lagrangian framework was developed at Politecnico di Milano and is used for the simulation of clouds containing supercooled water droplets (Bellosta et al. 2019). The Lagrangian framework permits the modeling of supercooled water droplets at an individual particle level. Subsequently, fundamental physics such as particle splashing on impact can be modeled. Techniques to allow for particle tracking in mesh with arbitrary motion and across non-conformal boundary interfaces are used within this work to simulate clouds entrained within rotorcraft flow fields (Morelli et al. 2019).

Finally, an icing solver is utilized to predict the rate of ice accretion. The in-house code PoliMIce is used for computing the ice accretion (Gori et al. 2015a). The PoliMIce software library provides state-of-the-art ice formation models, including the Myers model (Myers 2001), the modified Myers model (Gori et al. 2015b), and the local exact solution of the Stefan problem (Gori et al. 2017). Myers introduced a rime limit thickness,  $B_g$ , as a criterion for the selection of the ice regime. The rime ice limit thickness describes the condition at which the glaze regime can first appear. If the ice thickness  $B < B_g$  or  $B_g < 0$ : the *rime* ice accretion law is used, while if the ice thickness  $B > B_g$ : the *glaze* ice

accretion law is used. The rate of rime ice accretion can then be written as

$$\frac{\partial B}{\partial t} = \frac{\beta LWC U_\infty}{\rho_{\text{rime}}}, \quad (1)$$

while the rate of glaze ice accretion reads

$$\frac{\partial B}{\partial t} = \frac{1}{\rho_{\text{glaze}} L_F} (\dot{Q}_{\text{down}} + \dot{Q}_{\text{up}}), \quad (2)$$

where the collection efficiency, droplet liquid water content, and freestream velocity are, respectively, denoted by  $\beta$ , LWC, and  $U_\infty$ . Here,  $\rho$  represents the density of ice which depends on the ice regime.  $\dot{Q}_{\text{down}}$  and  $\dot{Q}_{\text{up}}$  are the heat fluxes exchanged on the ice/water phase change interface.  $L_F$  is the latent heat required for melting ice.

The model used in this work to capture the complex experimental ice shapes is the local exact solution of the unsteady Stefan problem for the temperature profiles within the ice layer in glaze conditions (Gori et al. 2017). The ice shapes are then computed using a multi-step approach, whereby non-linear ice accretion is accounted for by iteratively updating the surface solution on which the ice accretes. The efficacy of the approach for predicting ice shapes on two-dimensional oscillating airfoils is highlighted in Morelli et al. (2020).

## 2.2 Turbulent flow simulation and noise prediction

SU2 has been developed with the task of solving partial differential equations (PDE) and constrained optimization problems on general unstructured meshes. The core of the suite is a Reynolds-averaged Navier–Stokes (RANS) solver capable of simulating the compressible, turbulent flows that are representative of many problems in aerospace and mechanical engineering. The finite volume method (FVM) is applied on arbitrary unstructured meshes using a standard edge-based data structure on a dual grid with control volumes constructed using a median-dual, vertex-based scheme. Several numerical fluxes like Jameson–Schmidt–Turkel (JST), Upwind Roe, and variants of the AUSM schemes are implemented and slope limiters enable second-order space integration. For time discretization SU2 uses a second-order dual-time stepping method with an outer time loop to march through the physical time and of an inner loop which is usually a pseudo-time iteration or a (quasi-)Newton scheme. An Arbitrary Lagrangian–Eulerian (ALE) formulation of the RANS equations is implemented in SU2 to account for unsteady grid motion, in which the convective fluxes are adjusted to obtain solutions on arbitrarily moving grids.

In the SU2 suite, of particular relevance to this work are the scale-resolving and acoustic prediction capabilities. For

scale-resolving simulation, the Enhanced Delayed Detached Eddy Simulation (EDDES) based on the Spalart–Allmaras (SA) turbulence model with a shear-layer adapted (SLA) sub-grid scale model (Shur et al. 2015) was used. In addition, to limit the numerical dissipation in the LES part of the EDDES solver, the inviscid flux is computed using the Simple Low Dissipation Advection Upstream (SLAU2) scheme (Kitamura and Hashimoto 2016). The EDDES-SA solver has been demonstrated to successfully predict separated flows (Molina 2015; Molina et al. 2019), especially those induced by icing on the wing leading edge (Molina et al. 2020)

For aeroacoustic prediction, a solid-surface Ffowcs Williams and Hawkings (FWH) analogy is used. To that end, Farassat’s Formulation-1A (F1A) (Farassat 2007) has been implemented in the SU2 suite (Zhou et al. 2017; Icke et al. 2020; Morelli et al. 2022) which computes far-field noise from surface pressure fluctuations extracted from a preceding unsteady EDDES-SA simulation.

### 2.3 Machine learning using Bayesian neural networks

In this section, we first introduce the classical NN (Rosenblatt 1958), followed by the BNN (Bate et al. 1998) that captures the uncertainty of NN model parameters and finally the pSVGd (Chen and Ghattas 2020) algorithm used to numerically constructs a BNN.

#### 2.3.1 Neural networks

A NN is a function  $f$  mapping input feature  $x$  to output target variable  $y$ . We write  $\hat{y} = f(x)$  where the hat denotes the predicted value from the NN. For illustration, we focus solely on densely connected feedforward NNs (also known as multi-layer perceptrons), although more advanced NN architectures (e.g., convolutional and recurrent NNs) may also be adopted within our overall ML framework. The hyperparameters that define the NN architecture (e.g., number of layers, nodes per layer, and activation functions) are selected prior to NN training, and good hyperparameter

choices can be sought systematically through validation or cross-validation. Once the hyperparameters are chosen, the NN training then focuses on determining the weight and bias parameters of all NN layers. We denote the collection of all such trainable parameters by  $\omega$ . Then, more explicitly, we write the prediction under a particular NN parameter setting as  $\hat{y} = f(x; \omega)$


Given  $N$  training points in the form of input–output pairs  $(x_T, y_T) = \{x_n, y_n\}_{n=1}^N$ , NN training seeks  $\omega^*$  that minimizes a loss function reflecting the degree of mismatch between NN predictions and the true target values for these training points. In regression problems, for example, the mean squared error (MSE) loss is often used:

$$\omega^* = \arg \min_{\omega} \left\{ \frac{1}{N} \sum_{n=1}^N [f(x_n; \omega) - y_n]^2 \right\} \dots \tag{3}$$

This optimization problem is typically solved via gradient-based algorithms, such as stochastic gradient descent (LeCun et al. 2012; Robbins and Monro 1951) or ADAM (Kingma and Ba 2015). In any case, the solution to  $\omega^*$  is single valued, and any new prediction made by the trained NN at a new input  $\hat{y}_{\text{new}} = f(x_{\text{new}}; \omega^*)$  would also be single valued. No uncertainty information is produced to convey the confidence or quality of these estimates that would be affected by, for example, how many points were used to train the NN and how noisy those training points were. Having uncertainty quantification for the model and its predictions will be particular important to support safety- and mission-critical decision-making, such as in icing detection.

#### 2.3.2 Bayesian neural networks

We adopt a probabilistic approach that seeks to produce not just a single-valued prediction, but a *distribution* of plausible prediction values. This distinction is demonstrated in Fig. 3. In particular, we take a Bayesian approach and treat  $\omega$  as a random variable with an associated probability density function (PDF) that reflects the uncertainty



**Fig. 3** Illustration of a simple NN and BNN, each with an input layer, two hidden layers, and an output layer. Trainable parameters are depicted in red as deterministic values in the NN and as probability distributions in the BNN

in its values. When new training data become available, the PDF and therefore uncertainty of  $\omega$  is updated via Bayes' rule:

$$p(\omega | x_T, y_T) = \frac{p(x_T, y_T | \omega)p(\omega)}{p(x_T, y_T)} = \frac{p(y_T | x_T, \omega)p(\omega)}{p(y_T | x_T)}, \quad (4)$$

where  $p(\omega | x_T, y_T)$  is the posterior PDF on the weight parameters;  $p(\omega)$  is the prior PDF;  $p(y_T | x_T, \omega)$  is the likelihood PDF; and  $p(y_T | x_T)$  is a normalizing term independent of  $\omega$  often known as the model evidence or marginal likelihood. Solving the Bayesian inference problem then constitutes computing or characterizing the posterior  $p(\omega | x_T, y_T)$  for a given training dataset  $(x_T, y_T)$ . Once the posterior is available, we can propagate its uncertainty to predictions via Monte Carlo sampling  $\omega^{(i)} \sim p(\omega | x_T, y_T)$  and correspondingly produce samples of  $\hat{y}_{\text{new}}^{(i)} = f(x_{\text{new}}; \omega^{(i)})$  that represent the posterior predictive uncertainty—that is, a distribution of plausible prediction values—at any new input  $x_{\text{new}}$ .

Solving for the posterior is extremely difficult for NNs, since  $\omega$  can easily reach thousands, even millions of dimension for large NNs intended to model complex scientific and engineering processes. While Markov chain Monte Carlo algorithms (Gilks et al. 1996; Andrieu et al. 2003; Robert and Casella 2004; Brooks et al. 2011) are typically used for Bayesian computations for their ability to generate samples from the exact posterior, these methods do not scale well to such high dimensions in practice. In contrast, variational inference (VI) (Blei et al. 2017; Zhang et al. 2019) formulates an optimization problem that seeks the best approximation to the posterior from within a class of parameterized distributions in lieu of sampling the posterior. If we use  $q(\omega; \theta)$  to denote the approximate posterior with  $\theta$  representing the parameterization of these approximating distributions (e.g., if the  $q$  is from the family of all Gaussian distributions, then  $\theta$  is the mean  $\mu$  and covariance  $\Sigma$  of a Gaussian) and the VI formulation then seeks  $\theta^*$  that minimizes the Kullback–Leibler divergence (a measure of farness between two distributions) from the true posterior to the approximate posterior:

$$\theta^* = \arg \min_{\theta} D_{\text{KL}}[q(\omega; \theta) \| p(\omega | x_T, y_T)]. \quad (5)$$

Since an optimization task is only concerned with finding the single optimum point, it is more scalable to higher dimensions than a sampling task that needs to portray the entire landscape of the posterior distribution. Hence, by converting a sampling task to an optimization one, VI effectively trades off accuracy of posterior representation for scalability, making it more suitable for BNNs.

### 2.3.3 Projected stein variational gradient descent

We focus on two particle-based VI methods: Stein variational gradient descent (SVGD) (Liu et al. 2016) and projected SVGD (pSVGD) (Chen and Ghattas 2020). With particle representations, these methods are able to capture correlation and non-Gaussian structures of distributions that the commonly used mean-field Gaussian VI cannot.

SVGD is derived by leveraging the relationship between the (functional) gradient of objective in Eq. (5) to the Stein discrepancy, the latter which can be approximated using a set of particles (see Liu et al. 2016 for a detailed derivation). A gradient descent procedure can then be formed to iteratively minimize this gradient, as summarized by

$$\omega_i^{\ell+1} \leftarrow \omega_i^{\ell} + \epsilon_{\ell} \hat{\phi}^*(\omega_i^{\ell}), \quad (6)$$

where  $\omega_i^{\ell}$ ,  $i = 1, \dots, m$  is the location of the  $i$ th particle at optimization iteration  $\ell$ ,  $\epsilon_{\ell}$  is the learning rate that can also be adapted with  $\ell$ , and  $\hat{\phi}^*(\omega)$  is the update direction. In turn, the update direction can be expanded as

$$\hat{\phi}^*(\omega) = \frac{1}{M} \sum_{m=1}^M \left[ k(\omega_m^{\ell}, \omega) \nabla_{\omega_m^{\ell}} \log p(\omega_m^{\ell} | x_T, y_T) + \nabla_{\omega_m^{\ell}} k(\omega_m^{\ell}, \omega) \right] \quad (7)$$

with  $k(\cdot, \cdot)$  being a positive definite kernel. Furthermore, the gradient of the true log-posterior in the above equation can be evaluated via the sum of gradients of log-likelihood and log-prior, since the gradient of the log model evidence with respect to  $\omega$  is zero. The overall effect is an iterative transport of a set of particles to best match the target posterior distribution  $p(\omega | x_T, y_T)$ . One observed drawback to SVGD, however, is the tendency for particles to collapse toward the mode of the distribution as the dimensionality becomes high relative to the number of SVGD particles, thereby (potentially drastically) under-representing the uncertainty (Chen and Ghattas 2020).

pSVGD (Chen and Ghattas 2020) aims to address this issue by projecting parameters into a low-dimensional subspace and conducting the particle-based inference in this subspace. This allows inference to take place with a much higher particle-to-dimension ratio, which can mitigate the posterior collapsing phenomenon from SVGD. The specific space selected for projection is a likelihood-informed subspace (LIS). Components with the greatest likelihood-to-prior ratio are identified for inference while components with a negligible likelihood-to-prior ratio are assumed to be unmodified from the prior and need not be inferred directly. In other words, we will focus and work in only the subspace where the data can make a significant update from prior to posterior. We refer readers to Chen

and Ghattas (2020) for a detailed derivation and present a brief description below.

The LIS is determined by first calculating the average (log)-likelihood Hessian  $-\mathbb{E}_\omega[\nabla_\omega^2 \mathcal{L}]$  where  $\mathcal{L} = \log p(x_T, y_T | \omega)$ .  $-\nabla_\omega^2 \mathcal{L}$  can be calculated either directly via back-propagation of the NN to evaluate the second derivative or much more inexpensively via the Gauss-Newton approximation  $-\nabla_\omega^2 \mathcal{L} \approx -(\nabla_\omega \mathcal{L} \nabla_\omega \mathcal{L}^T)$  needing only first-derivative information. Next, we solve the generalized eigen-decomposition problem:

$$-\mathbb{E}_\omega[\nabla_\omega^2 \mathcal{L}] \psi_i = \lambda_i \Gamma_0^{-1} \psi_i, \quad i = 1, \dots, r, \tag{8}$$

where  $\lambda_i$  are the  $r$  leading eigenvalues that represent the relative influence of the likelihood over the prior,  $\psi_i$  are the corresponding eigenvectors, and  $\Gamma_0^{-1}$  is the inverse of the prior covariance. Therefore, eigenvalues below a certain threshold correspond to components with negligible influence from the prior, leaving the prior effectively unchanged. These negligible eigenmodes can then be truncated, and the LIS is constructed with the surviving eigenvectors corresponding to non-negligible eigenvalues. The choice of threshold value for truncation is application specific. For lower dimensional, mostly linear models, a threshold value  $\mathcal{O}(1)$  is sufficient. For larger, more complex models, a lower threshold value may be necessary. The original authors of pSVGd (Chen and Ghattas 2020) use a value of 0.01.

With the LIS found, the prior parameters  $\omega_0$  can be decomposed into the projected component  $\omega_0^r$  and the its complement  $\omega_0^T$  (subscript 0 indicates that the decomposition is done under the prior distribution). The complement  $\omega_0^T$  is then frozen at the prior projection. The projected component of the parameters will be inferred under the low-dimensional projected variable  $\chi$  using the same iterative procedure in SVGd, where  $\omega^r = \Psi^r \chi$ . The original full-parameter variable after pSVGd inference can then be reconstructed at the end:

$$\omega^\ell = \omega_0^T + \Psi^r \chi^\ell. \tag{9}$$

Performing Bayesian inference on  $\chi$  instead of  $\omega$  permits a much lower-dimensional problem. There is also secondary

benefit that by truncating to the  $r$  components most informed by data, the complement components that would otherwise be more amenable to fitting to data noise are now regularized, further improving predictive performance.


### 3 Results

#### 3.1 Ice prediction


The PoliMice (Gori et al. 2015a) tool kit is utilized for generating a large dataset of ice shapes which can subsequently be harnessed by the NN model. The icing conditions selected are based on the experimental work from Shin and Bond (Shin and Bond 1992). They conducted icing experiments in the Icing Research Tunnel (IRT) at the NASA Lewis Research Center. Their work studied the repeatability of ice shapes on a model wing at a range of temperatures varying from 247 to 271K. This work presents the numerical predictions of four ice shapes at different temperatures during the high-speed conditions for icing validation. The icing conditions selected are summarized in Table 1, while the ice predictions are presented in Fig. 4. The model used in the icing experiments was a straight wing with NACA0012 profile. The ice accretion simulations are consequently assumed to be two-dimensional due to the wing model being straight. The wing is set at an angle of attack of 4°. The duration which the wing is exposed to icing conditions last for 420 s. The wind tunnel airspeed was approximately  $M = 0.3$ . As previously mentioned, a range of different temperatures are investigated. The Liquid Water Content (LWC) and Mean Volume Diameter (MVD) of the supercooled water droplets is, respectively,  $0.55 \text{ g/m}^3$  and  $20 \text{ }\mu\text{m}$ . The numerical prediction shown in Fig. 4 are in relatively good agreement with the experimental measurements. The ice predictions at the lower temperature conditions displayed in Fig. 4a and b closely resemble the ice measurements. There are slightly greater discrepancies when moving to the higher temperatures as shown in Fig. 4c and d. The numerical predictions tend to predict the presence of glaze ice earlier than the experiments.

**Table 1** Icing conditions based on the experimental work at the NASA Lewis Icing Research Center (Shin and Bond 1992)


Run ID	Accretion time (s)	Atmosph. pressure (Pa)	Freestream velocity (m/s)	Outside air temperature (K)	Liquid water content ( $\text{g/m}^3$ )	Mean volume diameter ( $\mu\text{m}$ )
Rime 1	420	1,01,325	102.82	247.0	0.55	20
Rime 2	420	1,01,325	102.82	262.0	0.55	20
Glaze 1	420	1,01,325	102.82	267.5	0.55	20
Glaze 2	420	1,01,325	102.82	271.0	0.55	20




(a) Rime ice regime 1 at 247.0K.



(b) Rime ice regime 2 at 262.0K.



(c) Glaze ice regime 1 at 267.5K.



(d) Glaze ice regime 2 at 271.0K.

**Fig. 4** Comparison of ice predictions with measurements at different temperatures. The ice shape is displayed in non-dimensional form where the coordinates,  $x$  and  $y$ , are normalized by the wing chord length,  $c$

**Table 2** Test matrix to produce sample of rime and glaze ice shapes

Accre- tion time (s)	Atmos- pheric pressure (Pa)	Freestream velocity(m/s)	Outside air tem- perature (K)	Liquid water content (g/m <sup>3</sup> )	Mean volume diameter (μm)
420	1,01,325	102.82	246 <sup>(+0.5)</sup> 271	0.55	20, 40

Proceeding the icing validation assessment, the subsequent goal of this work was to generate a dataset of ice shapes for the NN model. The experimental work from Shin and Bond was selected with this in mind. In their work they studied the effect of 7 different temperatures on the ice shape. Due to the NN model requiring a significantly larger dataset, this work increments the temperature by 0.5K to provide 50 samples. Additionally, ice shapes at a MVD of 40 μm are simulated to provide 100 samples in total. The icing conditions used for all the samples are summarized in Table 2.



### 3.2 Mesh generation

The focus on modeling small-scale vortical structures using scale-resolving simulations capabilities in SU2 requires three-dimensional grids. An in-house grid generation software called uhMesh (Dussin et al. 2009) is utilized to automatically generate a three-dimensional mesh from a two-dimensional ice shape. This is achieved by extruding the ice shape for 1 chord length in the spanwise direction. The assumption is that while the ice geometry contains no variation in the spanwise direction, the flow does due the scale-resolving modeling approach. Simulating a fully three-dimensional icing dataset of this size would be beyond the current hardware capabilities. Additionally, the ice prediction approach based on Myers' model (Myers 2001) would not capture the spanwise variation in the ice shape. To capture such three-dimensional roughness in the ice shape would require an approach, such as the one suggested by Szilder and Lozowski (2018). However this would be highly challenging to mesh in order to determine the aerodynamics and aeroacoustics of the iced wing.


Figure 5 illustrates the mesh generation for the clean wing using this approach. With scale-resolving simulations in mind, a boundary layer with a height of approximately  $0.5c$  is used to generate a structured region surrounding the wing. The boundary layer at the wall is resolved to ensure  $y^+ < 1$  everywhere. The mesh is extruded in the spanwise direction for 1 chord length and discretized by 150 elements. The total number of elements for each mesh is approximately 16 M

with it varying by as much as 500k depending on the ice shape.


### 3.3 Turbulent flow simulation and noise prediction for different ice shapes

To compare the turbulent flow field and far-field noise features of the clean and iced wings, the clean NACA0012 airfoil as well as a rime- and a glaze-iced airfoils associated with Rime 1 and Glaze 2 in Table 1 are presented. The scale-resolving EDDES-FWH approach outlined in Sect. 2.2 is used to model the fine-scale turbulent structures induced by the ice accretion. With ice detection on rotorcraft in mind, the  $1c$  section of the wing is subsequently set to pitch around its quarter-chord to represent the cyclic pitching motion of a rotor blade in forward flight. While this is a primitive way to model a rotor blade in forward flight, it was a requirement due to rotorcraft icing simulations requiring fully three-dimensional simulations which are computationally very costly. The sinusoidal pitching motion of the wing can be described by  $\alpha = 4^\circ \pm 4^\circ$  and pitching with a frequency of 6.0 Hz. The mean angle of attack is consistent with the icing experiments presented in Sect. 3.1. A time step of  $\Delta t = 0.05c/U_\infty = 0.00025$  s is used, which results in 666 time steps within each pitching cycle. The flow statistics and acoustics are computed after the 5th pitching cycle, over 10 pitching cycles.

The instantaneous flow fields at  $\text{AoA} = 4.0^\circ$  on the upstroke are visualized by the iso-surfaces of the Q-criterion




(a) Close-up of the mesh surrounding the wing.



(b) Isometric view of the surface mesh.

Fig. 5 Visualization of automated 3D mesh



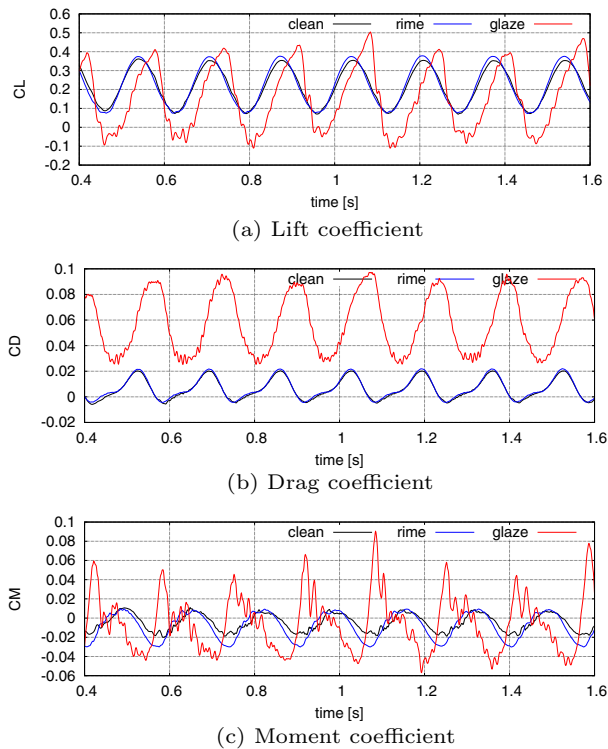
**Fig. 6** Q-criterion iso-surface for the clean and iced wings colored by non-dimensionalized velocity magnitude,  $U/U_\infty$

colored by non-dimensionalized velocity magnitude, as shown in Fig. 6. The effect of leading-edge icing is evident. While the clean and rime ice airfoil show mostly attached flows over the suction side, detached eddies signifying separated flow can clearly be seen in the more severely iced (glaze ice) airfoil, where the Q-criterion shows more fine-scale, three-dimensional turbulent structures over the entire suction side. Note that even though the ice shapes are quasi-2D, the resolved turbulent structures by EDDES are very much three-dimensional, especially in the glaze ice case where the Kelvin–Helmholtz instability is triggered immediately past the ice horn and develops quickly into 3D turbulence.

The time histories of the lift, drag, and moment coefficients of the rime and glaze ice wings, computed using the EDDES-SA simulations are compared in Fig. 7. In the case of rime ice, the force and moment coefficients are

very similar to those of the clean airfoil. Small differences between the clean and rime ice only occur at the peaks of the values which are associated with the maximum and minimum angle of attack. On the contrary, the glaze ice airfoil exhibits significantly deteriorated lift and strongly elevated drag, signifying flow separation over a large portion of the airfoil suction side. The moment coefficient is also severely effected during the glaze ice regime.

This is confirmed by examining the wall shear stress of the three cases shown in Fig. 8. Both clean and rime ice airfoils show trailing edge separation due to the dynamic effect of the upstroke. Surface noise footprints visualized by the instantaneous wall pressure fluctuation ( $p' = p - \bar{p}$ ) of the clean, rime, and glaze ice airfoil sections are compared in Fig. 9. The clean and rime ice airfoils exhibit similar surface noise footprints. On the other hand, the glaze airfoil exhibits a highly irregular noise footprint pattern



**Fig. 7** Time histories of lift, drag, and moment coefficients of the clean, rime, and glaze ice airfoils computing using the EDDES-SA simulations

with significant spatial frequency content (Fig. 10). One may already speculate that the far-field noise signature of the glaze ice should be associated with elevated high-frequency contents. This is confirmed by the far-field noise spectra shown in Fig. 11—while the rime ice airfoil shows a slight increase of high-frequency noise over the clean airfoil, the glaze airfoil has a significantly elevated noise level over the entire high-frequency range up to the cut-off. Clearly, the more severe the leading-edge ice accretion is, the more elevated the high-frequency range of the far-field noise tends to be. This can be exploited as a distinguishing feature in a machine learning model to build a mapping between the far-field noise level and ice shape as well as the associated aerodynamic performance of the ice airfoil.


### 3.4 Machine learning predictions of aerodynamic performance under icing conditions

Following Sect. 2.3, we build BNNs to predict aerodynamic performance metrics from inputs of acoustic measurements. Specifically, the BNN takes input of a compressed principal

component analysis (PCA) representation of the power spectral density (PSD) of the acoustics time series and predict the minimum, mean, and maximum values of the lift, drag, and moment coefficients over the duration of the time series window. To achieve this, we pre-process the acoustic time series in the training dataset by computing the log values of PSD for each data signal covering the frequency range 17–1000 Hz and perform a PCA to find and truncate to only the 6 leading components. All PSDs are then projected into this lower-dimensional representation. All inputs are then standardized using the mean and standard deviation statistics from across all training data simulations. For the output variables, we extract from each simulation’s aerodynamic time series the minimum, maximum, and mean over the time window for  $C_L$ ,  $C_D$ , and  $C_M$ , a total of 9 quantities of interest (QoIs). Overall, the BNN maps from a 6-dimensional input to a 9-dimensional output. The precise architecture is shown in Table 3, with one hidden layer of 15 neurons leading to a total of 249 model parameters. We purposefully favor a small BNN architecture due to the relatively few training samples available.

To set up the BNN, the prior  $p(w)$  from Eq. (4) is chosen to be independent Gaussian  $w_k \sim \mathcal{N}(0, 0.2^2)$ , and the likelihood  $p(y_T | x_T, w)$  is chosen to reflect an additive independent Gaussian data noise:  $y_k = f(x_k; w) + \epsilon_k$ ,  $\epsilon_k \sim \mathcal{N}(0, 0.2^2)$ . We use 600 particles to perform SVGD and pSVGd for our 249-parameter BNN. In order to obtain a robust measure of predictive performance of the BNN, we perform fivefold testing as follows. We divide the overall dataset of 93 simulations into fivefolds of 18–19 simulations each. Then, 5 BNN models are built where each model is trained on 4 of these folds and tested on the remainder fold. This allows us to obtain test performance statistics where each of the 93 data points gets a chance to play the role of a testing point.

Figure 12 presents both the SVGD and pSVGd BNN predictions, with each panel showing an aggregated result from the 5 testing scenarios. Using BNNs, we now have uncertainty information along with predictions. These uncertainties are depicted as error bars ( $\pm 1$  standard deviation) in the plot, which combine both the epistemic uncertainty calculated via SVGD/pSVGd and the output aleatoric uncertainty specified in the likelihood above (i.e., the additive Gaussian data noise  $\epsilon_k$  with  $\sigma_\epsilon = 0.2$ ). Note that input aleatoric uncertainty (i.e., measurement error of the acoustic signal) is not considered in this work, and requires further investigation since they are often not additive independent Gaussian. The SVGD and pSVGd results match well overall. pSVGd recovers slightly greater uncertainty of around 16% across



**Fig. 8** Skin friction coefficient,  $C_{F_x}$ , for the clean and iced wings

all predictions compared to SVGD which is amenable to underpredicting uncertainty due to particle collapse. The reduced dimension from pSVG D is about 80% lower than that of SVGD. While pSVG D uses additional approximation in its eigenvalue truncation, we believe this truncation error is outweighed by the benefit from a greater particle-to-dimension ratio resulting from the dimension reduction, leading to an overall more accurate posterior representation for pSVG D.

Posterior predictive expectation of the absolute and normalized root-mean-square error (RMSE) (i.e., posterior-averaged RMSE) for each QoI by the pSVG D model are shown in Table 4, showing good agreement with the

true target values. However, we draw caution to interpreting these results, since BNNs are not meant to minimize error directly (error is not its loss function) but rather to accurately capture the uncertainty in the NN as justified by the available training data at hand. As such, while it is still valuable to understand the error performance from these BNNs, one must be careful to not judge these models based on error alone.

In general, we also observe rime ice predictions to achieve lower error and lower uncertainty than glaze ice predictions, which matches our intuition where generally more complex aerodynamic flows are caused by the more severe glaze icing conditions. These trends are shown in

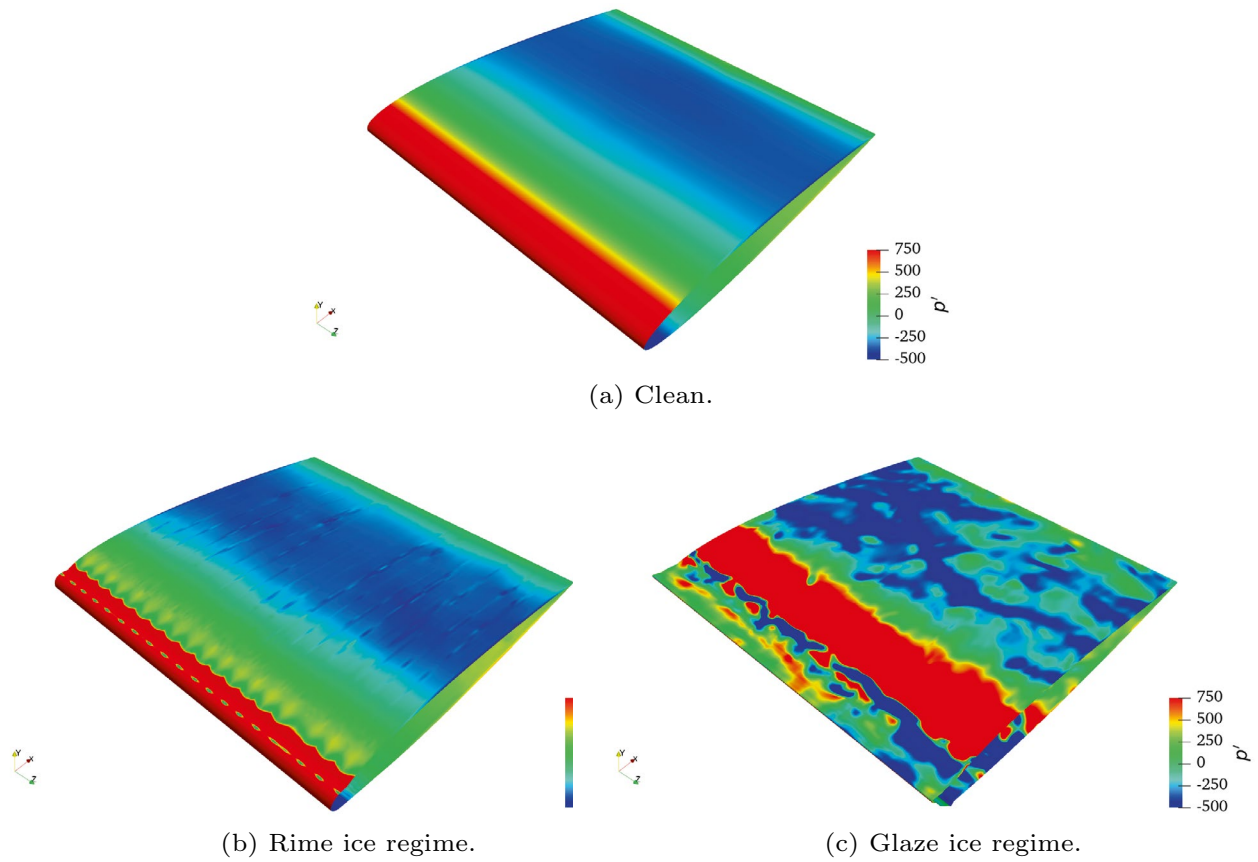


Fig. 9 Instantaneous pressure fluctuations,  $p'$ , for the clean and iced wings




Fig. 10 Time history of the pressure fluctuations,  $p'$ , for the clean and iced wings

Fig. 13 where we gather histograms of the prediction errors and uncertainty across all testing scenarios. Figure 14 further presents the posterior predictive distributions for  $C_D$ ,  $C_L$ , and  $C_M$  QoIs between a specific rime ice case and a specific glaze ice case, illustrating an overall good agreement with the predictive distributions and the true target values, and the greater uncertainty in predicting under glaze than rime ice.




Fig. 11 Far-field sound spectra of the clean NACA0012, rime, and glaze ice airfoil sections computed  $20c$  below the center of rotation of the pitching airfoil

**Table 3** BNN architecture

Layer	Name	Type	No. of neurons	Activation Fcn
1	Input	Dense	6	–
2	Hidden	Dense	15	Swish
3	Output	Dense	9	Linear

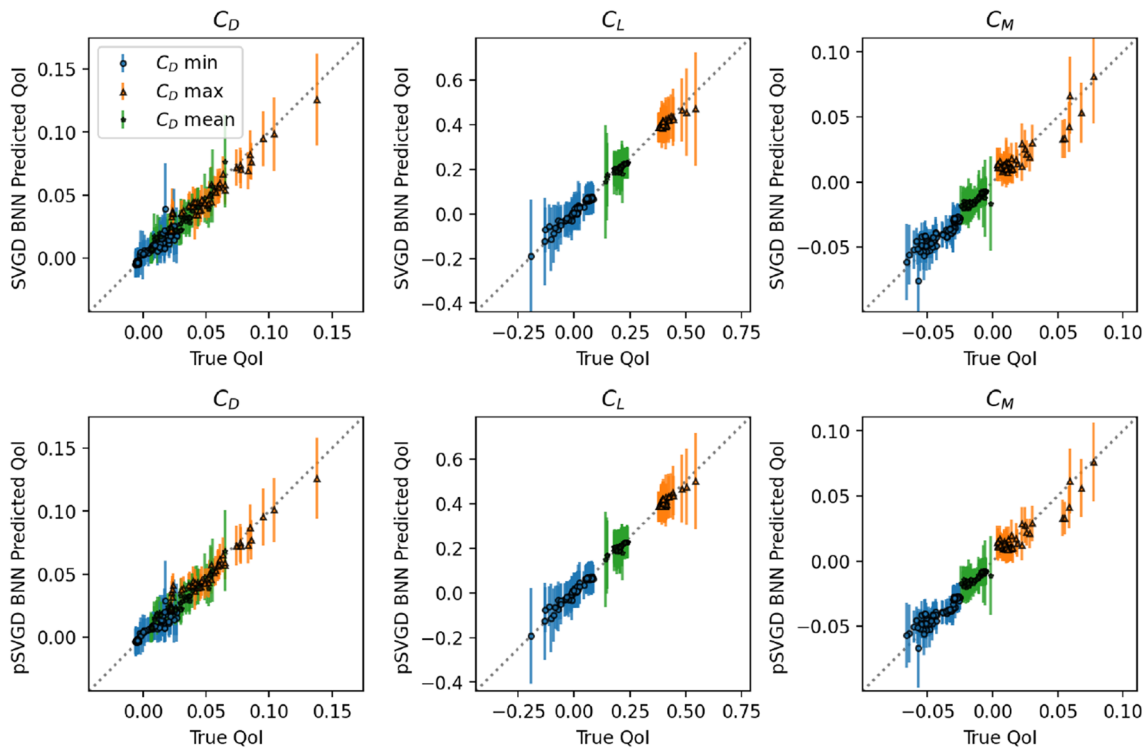
### 4 Conclusion

This work introduces an original methodology for developing a Bayesian neural network acoustic ice detection system. The acoustic signatures generated by different iced wings are harnessed to predict their aerodynamic performance. An inter-disciplinary methodology is adopted to develop the Bayesian neural network acoustic ice detection system.

**Table 4** The posterior predictive expectations of the normalized and absolute RMSE for the pSVG D model (i.e., posterior-averaged RMSE)


QoI	Normalized error	Absolute error
$C_D$ min	0.16	0.0044
$C_D$ max	0.08	0.0052
$C_D$ mean	0.20	0.0037
$C_L$ min	0.20	0.0140
$C_L$ max	0.08	0.0142
$C_L$ mean	0.21	0.0104
$C_M$ min	0.14	0.0048
$C_M$ max	0.06	0.0058
$C_M$ mean	0.11	0.0030

Normalized errors are calculated based on transforming the true test target values and predictions by the mean and standard deviation of the true training target values, which in turn can be interpreted roughly as percent error



**Fig. 12** Predicted versus true values for all QoIs, using SVGD (top row) and pSVG D (bottom row). The two sets of results overall match well with each other. More specifically, pSVG D recovered approximately 16% more of the Bayesian predictive uncertainty relative to

its SVGD counterpart that is amenable to underpredicting uncertainty due to particle collapse. pSVG D also achieved 80% dimension reduction of the parameter space compared to SVGD

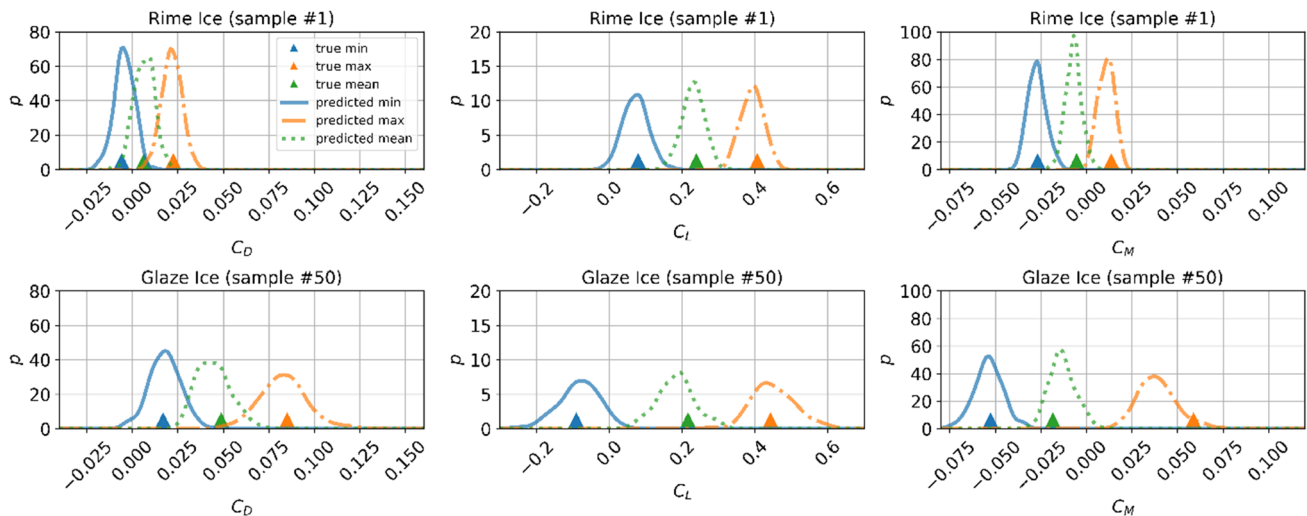


**Fig. 13** MSE of the posterior predictive distributions (left) and standard deviation of the posterior predictive distributions (right) averaged over all QoIs. The predicted QoIs are normalized to allow direct comparison and averaging

The PoliMIce toolkit is utilized to predict a test matrix of iced airfoils. The aerodynamics and aeroacoustics associated with the iced airfoils are, respectively, solved in SU2 using high-fidelity EDDES scale-resolving simulations and the FWH-F1A approach. A Bayesian neural network is then trained using projected Stein variational gradient descent to provide (1) mapping from observed far-field broadband

noise levels to predictions of aerodynamic performance indicators and (2) uncertainty information reflecting the neural network’s confidence in its predictions. While the icing simulation, aeroacoustic prediction as well as the construction of Bayesian neural network model all require potentially extensive computational resources in the *offline* phase, the prediction of aerodynamic performance indicators based on measured far-field noise level can be obtained rapidly in the time-critical *online* phase in-flight.

To the best of our knowledge, this work is the first to employ high-fidelity aeroacoustic simulations to develop correlations between aerodynamic performance and far-field noise level of iced airfoil sections via machine learning techniques. That being said, it should also be noted that the current work as-is only serves as a proof-of-concept study. It is clear that the pitching wing considered in this study is a highly simplified representation of the complex rotorcraft icing problem. Various noise sources such as those from tail rotors and engines, as well as rotor–fuselage interaction must be taken into account and differentiated in developing an ice detection system using the multi-disciplinary framework proposed in the current work.



**Fig. 14** Comparison of posterior predictive distributions for a test sample of rime ice (top row) and glaze ice (bottom row) for  $C_D$  (left),  $C_L$  (center), and  $C_M$  (right). Ground truth target values are indicated with triangles

**Acknowledgements** The authors gratefully acknowledge the high-performance computing resources provided by the UK National Supercomputing Service via ARCHER2 (<http://www.archer2.ac.uk>), on Project e734 (High-Fidelity Aerodynamic and Aeroacoustic Simulations for Urban Air Mobility (UAM) Concepts). JH and XH are supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Number DE-SC0021397. This paper was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, make any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

## Declarations

**Conflict of interest** On behalf of all authors, the corresponding author states that there is no conflict of interest.

**Replication of results** The numerical results presented in this manuscript can be reproduced using the input file and the codes accessible at this link: [https://uob-my.sharepoint.com/:f:/r/personal/jb21343\\_bristol\\_ac\\_uk/Documents/SAMO\\_SI\\_Data?csf=1&web=1&e=RJPUI1](https://uob-my.sharepoint.com/:f:/r/personal/jb21343_bristol_ac_uk/Documents/SAMO_SI_Data?csf=1&web=1&e=RJPUI1).

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- Andrieu C, de Freitas N, Doucet A, Jordan MI (2003) An introduction to MCMC for machine learning. *Mach Learn* 50:5–43
- Bassey CE, Simpson GR (2007) Aircraft ice detection using time domain reflectometry with coplanar sensors. In: 2007 IEEE aerospace conference. IEEE, pp 1–6. <https://doi.org/10.1109/AERO.2007.352857>
- Bate A, Lindquist M, Edwards IR, Olsson S, Orre R, Lansner A, De Freitas RM (1998) A Bayesian neural network method for adverse drug reaction signal generation. *Eur J Clin Pharmacol* 54:315–321
- Bellosta T, Parma G, Guardone A (2019) A robust 3D particle tracking solver for in-flight ice accretion using arbitrary precision arithmetic. In: VIII international conference on computational methods for coupled problems in science and engineering
- Blei DM, Kucukelbir A, McAuliffe JD (2017) Variational inference: a review for statisticians. *J Am Stat Assoc* 112(518):859–877. <https://doi.org/10.1080/01621459.2017.1285773>
- Bragg M, Gregorek G, Lee J (1986) Airfoil aerodynamics in icing conditions. *J Aircr* 23(1):76–81
- Bragg M, Basar T, Perkins W, Selig M, Voulgaris P, Melody J, Sarter N (2002) Smart icing systems for aircraft icing safety. In: 40th AIAA aerospace sciences meeting & exhibit, p 813
- Broeren AP, Bragg MB (2005) Effect of airfoil geometry on performance with simulated intercycle ice accretions. *J Aircr* 42(1):121–130
- Brooks S, Gelman A, Jones G, Meng X-L (eds) (2011) Handbook of Markov chain Monte Carlo. Chapman & Hall/CRC, Boca Raton. <https://doi.org/10.1201/b10905>
- Caliskan F, Hajiyev C (2013) A review of in-flight detection and identification of aircraft icing and reconfigurable control. *Prog Aerosp Sci* 60:12–34
- Chen P, Ghattas O (2020) Projected stein variational gradient descent. *Adv Neural Inf Process Syst* 33:1947–1958
- Chen X, Zhao Q, Barakos GN, Kusyumov A (2019) Numerical analysis of rotor aero-acoustic characteristics for ice detection. *Int J Aeroacoust* 18(6–7):596–620
- Cheng B, Han Y, Brentner KS, Palacios J, Morris PJ, Hanson D, Kinzel M (2018) Surface roughness effect on rotor broadband noise. *Int J Aeroacoust* 17(4–5):438–466
- Deiler C, Fezens N (2020) Performance-based ice detection methodology. *J Aircr* 57(2):209–223
- Dussin D, Fossati M, Guardone A, Vigeveno L (2009) Hybrid grid generation for two-dimensional high-reynolds flows. *Comput Fluids* 38(10):1863–1875
- Farassat F (2007) Derivation of formulations 1 and 1a of farassat. In: NASA/TM-2007-214853. NASA Langley Research Center
- Gent R, Dart N, Cansdale J (2000) Aircraft icing. *Philos Trans R Soc Lond Ser A* 358(1776):2873–2911. <https://doi.org/10.1098/rsta.2000.0689>
- Gilks WR, Richardson S, Spiegelhalter DJ (1996) Markov chain Monte Carlo in practice. Chapman & Hall, New York
- Gori G, Zocca M, Garabelli M, Guardone A, Quaranta G (2015a) PoliMIce: a simulation framework for three-dimensional ice accretion. *Appl Math Comput* 267:96–107. <https://doi.org/10.1016/j.amc.2015.05.081>
- Gori G, Zocca M, Guardone A (2015b) A model for in-flight ice accretion based on the exact solution of the unsteady stefan problem. In: 7th AIAA atmospheric and space environments conference, AIAA AVIATION, Dallas, p 3019
- Gori G, Parma G, Zocca M, Guardone A (2017) Local solution to the unsteady stefan problem for in-flight ice accretion modeling. *J Aircr* 55(1):251–262. <https://doi.org/10.2514/1.C034412>
- Hochreiter S, Schmidhuber J (1997) Long short-term memory. *Neural Comput* 9:1735–80. <https://doi.org/10.1162/neco.1997.9.8.1735>
- Hongerholt DD, Willms G, Rose JL (2002) Summary of results from an ultrasonic in-flight wing ice detection system. In: AIP conference proceedings, vol 615. American Institute of Physics, pp 1023–1028
- Icke RO, Baysal O, Moy A, Lopes L, Zhou BY, Diskin B (2020) Toward adjoint-based aeroacoustic optimization for propeller and rotorcraft applications. In: AVIATION 2020 Forum, AIAA-2020-3140
- Kingma DP, Ba JL (2015) Adam: a method for stochastic gradient descent. ICLR: International conference on learning representations. [arXiv:1412.6980v9](https://arxiv.org/abs/1412.6980v9)
- Kitamura K, Hashimoto A (2016) Reduced dissipation AUSM-family fluxes: HR-SLAU2 and HR-AUSM+-up for high resolution unsteady flow simulations. *Comput Fluids* 126:41–57
- Lecun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proc IEEE* 86(11):2278–2324. <https://doi.org/10.1109/5.726791>



- LeCun YA, Bottou L, Orr GB, Müller K-R (2012) Efficient BackProp. In: Montavon G, Orr GB, Müller K-R (eds) Neural networks: tricks of the trade. Springer, Berlin, pp 9–48. [https://doi.org/10.1007/978-3-642-35289-8\\_3](https://doi.org/10.1007/978-3-642-35289-8_3)
- Liu Q, Wang D (2016) Stein variational gradient descent: a general purpose Bayesian inference algorithm. In: Advances in neural information processing systems 29 (NIPS 2016), Barcelona, Spain, pp 2378–2386
- Melody JW, Başar T, Perkins WR, Voulgaris PG (2000) Parameter identification for inflight detection and characterization of aircraft icing. *Control Eng Pract* 8(9):985–1001
- Molina E (2015) Detached eddy simulation in SU2. Ph.d. thesis, Aeronautical Institute of Technology
- Molina EM, Zhou BY, Alonso JJ, Righi M, da Silva RG (2019) Flow and noise predictions around tandem cylinders using DDES approach with SU2. AIAA-2019-0326. <https://doi.org/10.2514/6.2019-0326>
- Molina ES, Silva DM, Broeren AP, Righi M, Alonso eY Juan J, Peng S-H, Schwamborn D, Revell A, Mockett C (2020) Application of DDES to Iced Airfoil in Stanford University Unstructured (SU2). In: Progress in hybrid RANS-LES modelling. Springer, Cham, pp 283–293
- Morelli M, Guardone A (2021) A simulation framework for rotorcraft ice accretion and shedding. *Aerosp Sci Technol* 129:107157
- Morelli M, Bellosta T, Guardone A (2019) Lagrangian particle tracking in sliding mesh for rotorcraft icing applications. In: 45th European rotorcraft forum. Warsaw, Poland, pp 1–13
- Morelli M, Zhou BY, Guardone A (2020) Acoustic characterization of glaze and rime ice structures on an oscillating airfoil via fully unsteady simulations. *J Am Helicopter Soc* 65(4):1–12
- Morelli M, Guardone A, Zhou BY (2022) Computational aeroacoustic analysis of a wing-tip mounted propeller and high-lift device. In: 28th AIAA/CEAS aeroacoustics 2022 conference, p 2939
- Myers TG (2001) Extension to the messinger model for aircraft icing. *AIAA J* 39(2):211–218. <https://doi.org/10.2514/2.1312>
- Padfield GD (2008) Helicopter flight dynamics: the theory and application of flying qualities and simulation modelling, 2nd edn. Blackwell Publishing, Oxford
- Robbins H, Monro S (1951) A stochastic approximation method. *Ann Math Stat* 22(3):400–407. <https://doi.org/10.1214/aoms/1177729586>
- Robert CP, Casella G (2004) Monte Carlo statistical methods. Springer, New York. <https://doi.org/10.1007/978-1-4757-4145-2>
- Rosenblatt F (1958) The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev* 65(6):386
- Roy S, Izad A, DeAnna RG, Mehregany M (1998) Smart ice detection systems based on resonant piezoelectric transducers. *Sens Actuators A* 69(3):243–250
- Schlegel T, Moser M, Zangl H (2015) Wireless and flexible ice detection on aircraft. Technical report, SAE Technical Paper
- Shin J, Bond TH (1992) Experimental and computational ice shapes and resulting drag increase for a naca 0012 airfoil. NASA Technical Memorandum No. 105743
- Shur ML, Spalart PR, Strelets MK, Travin AK (2015) An enhanced version of DES with rapid transition from RANS to LES in separated flows. *Flow Turbul Combust* 95(4):709–737
- Szilder K, Lozowski EP (2018) Comparing experimental ice accretions on a swept wing with 3d morphogenetic simulations. *J Aircr* 55(6):2546–2549. <https://doi.org/10.2514/1.C034879>
- Varadan VK, Varadan VV, Bao X-Q (1997) IDT, SAW, and MEMS Sensors for measuring deflection, acceleration, and ice detection of aircraft. In: Smart structures and materials 1997: smart electronics and MEMS, vol 3046. International Society for Optics and Photonics, pp 209–219
- Zhang C, Butepage J, Kjellstrom H, Mandt S (2019) Advances in variational inference. *IEEE Trans Pattern Anal Mach Intell* 41(8):2008–2026. <https://doi.org/10.1109/TPAMI.2018.2889774>
- Zhou BY, Albring T, Gauger NR, Ilario CR, Economou TD, Alonso JJ (2017) Reduction of airframe noise components using a discrete adjoint approach. AIAA-2017-3658. <https://doi.org/10.2514/6.2017-3658>
- Zhuge JC, Yu ZJ, Gao JS (2012) Ice detection based on near infrared image analysis. In: Applied mechanics and materials, vol 121. Trans Tech Publ, pp 3960–3964

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.