

WIND TURBINE

Stochastic Operator Residual Learning with Conformal Calibration for Risk-Aware Wind Turbine Maintenance

Pankaj Sharma

Department of Mechanical Engineering, Jagannath University, Jaipur, Rajasthan, India

Published October 2026

Abstract

Wind turbine maintenance decisions are frequently made under uncertainty, yet conventional degradation models often provide deterministic predictions that ignore the tail risks associated with unexpected component failure. We propose a stochastic operator residual learning framework with conformal calibration for risk-aware wind turbine maintenance. The proposed method learns a distribution over the residual field of governing fatigue and wear equations, rather than enforcing fixed physical laws through penalty terms, so that complex loading spectra and material heterogeneities can be captured without relying on potentially inaccurate constants. An ensemble of neural operators is then trained to separate epistemic from aleatoric uncertainty, and conformal prediction is applied to calibrate the resulting confidence intervals, yielding a remaining useful life distribution with guaranteed coverage. These calibrated distributions are subsequently passed to a risk-sensitive deep reinforcement learning agent whose objective combines power production gains, maintenance costs, and a conditional value-at-risk penalty on the remaining useful life. Furthermore, gradients from the policy update flow back through the calibration layer and the stochastic operator, which allows the uncertainty representation to be shaped by the risk profile required by the controller. The main contribution is a differentiable integration of stochastic residual learning, calibrated uncertainty quantification, and risk-constrained control within a single digital twin architecture. Simulation-based evaluation on representative turbine degradation scenarios will assess whether the framework reduces unexpected failures while preserving energy yield, and the results are expected to inform safer maintenance scheduling under uncertain operating conditions.

Keywords: wind turbine; predictive maintenance; neural operators; conformal prediction; risk-sensitive reinforcement learning; remaining useful life

1 Introduction

Wind turbine systems operate in highly stochastic environments where structural integrity is continuously compromised by complex fatigue and wear mechanisms driven by aerodynamic loads, mechanical vibrations, and environmental variability. Predictive maintenance has therefore become a central component of modern wind energy operations, as unplanned failures not only incur substantial repair costs but also reduce availability and increase the levelized cost of energy [1]. Traditional approaches rely heavily on physics-based models such as Miner's rule for cumulative damage [2] and finite element analysis for stress distribution [3]. However, these deterministic formulations often fail to capture the nonlinear dynamics of degradation under varying operating conditions, and they lack the real-time computational efficiency required for online decision-making.

Data-driven methods have gained prominence as a consequence, particularly deep learning architectures that learn degradation patterns directly from sensor data [4]. Nevertheless, purely black-box models frequently violate physical laws and generalize poorly outside their training distributions. Physics-informed neural networks (PINNs) address this limitation by embedding partial differential equations into the loss function [5], and they have been applied successfully to fluid-structure interactions and thermal dynamics [6]. Standard PINNs, however, struggle with high-dimensional PDEs and long-term temporal dependencies due to spectral bias and optimization difficulties. Neural operators, including the Fourier Neural Operator [7] and DeepONet [8], overcome several of these obstacles by learning mappings between infinite-dimensional function spaces, offering mesh-independent predictions and faster inference than traditional solvers.

Uncertainty quantification remains a critical challenge in safety-critical applications such as wind energy. Deterministic models cannot account for sensor noise or the epistemic uncertainty arising from limited training data. Bayesian deep learning provides a principled framework for capturing both aleatoric and epistemic uncertainties [9], with Monte Carlo dropout [10] and ensemble learning [11] serving as widely used approximations to posterior distributions. Conformal prediction offers finite-sample statistical guarantees for prediction intervals without assuming specific data distributions [12]. In degradation modeling, probabilistic state-space models such as hidden Markov models [13] and Gaussian processes [14] have been employed to track health indices over time. Coupling these uncertainty estimates directly with control actions, however, requires differentiable propagation of uncertainty through complex dynamical systems, an area where stochastic neural operators show considerable promise [15].

Decision-making in wind turbine operations involves balancing energy capture against component wear. Reinforcement learning has emerged as a powerful tool for adaptive control [16], [17], yet standard algorithms optimize expected rewards and can therefore permit catastrophic tail events. Risk-aware reinforcement learning incorporates measures of risk into the objective function, and conditional value-at-risk (CVaR) has become a coherent risk measure that focuses on the average loss in worst-case scenarios [18]. CVaR-constrained frameworks have been developed to ensure robustness against extreme outcomes [19]. Digital twins provide a foundational infrastructure for such integration, creating virtual replicas of physical assets synchronized via real-time data streams [20], and hybrid digital twins combine first-principles physics with data-driven corrections to improve fidelity [21]. High-fidelity aeroelastic simulators such as OpenFAST are often used as environments for training reinforcement learning agents [22], and surrogate models derived from these simulators accelerate policy learning [23].

The intersection of these fields—neural operators for PDE residuals, Bayesian uncertainty decomposition, conformal calibration, and CVaR-constrained reinforcement learning—represents a frontier in autonomous system management. While individual components exist in isolation, their integrated application to wind turbine fatigue prediction and control setpoint optimization under strict risk constraints remains largely unexplored. Previous studies have treated prediction and control separately, or have used simplified uncertainty models that do not propagate through the entire decision pipeline. The reliance on ensemble methods for uncertainty decomposition draws from broader machine learning practices [24], while differentiable propagation of probabilistic states aligns with research in probabilistic programming and gradient-based inference [25].

We propose a hybrid framework that couples stochastic neural operators with risk-sensitive deep reinforcement learning for predictive maintenance and performance optimization in wind turbine

systems. The core novelty lies in the use of a stochastic neural operator to map operating parameters directly to fatigue and wear PDE residual fields, thereby capturing complex degradation dynamics more accurately than traditional deterministic models. Ensemble methods and conformal calibration rigorously decompose epistemic and aleatoric uncertainties, and these calibrated distributions are propagated differentially through the degradation state estimation process to yield precise remaining useful life predictions. These probabilistic states then inform a physics-consistent reinforcement learning policy optimized under CVaR constraints, ensuring that maintenance scheduling and control actions remain robust against tail risks while maximizing energy capture. The key difference from existing methods is the end-to-end differentiability of the uncertainty representation, which allows the risk profile of the controller to shape the learned degradation model rather than treating prediction and control as independent modules.

The remainder of this paper reviews related work on neural operators, uncertainty quantification, and risk-aware reinforcement learning for wind energy systems; presents the preliminaries and problem formulation; details the proposed stochastic neural operator-coupled risk-sensitive reinforcement learning framework; describes the experimental setup and reports results and analysis; and discusses limitations and future directions and concludes the paper.

2 Related work

The proposed framework sits at the intersection of physics-informed degradation modeling, uncertainty quantification, and risk-sensitive control. Accordingly, we organize the discussion around three strands of literature that inform our design: neural operator learning for physical systems, probabilistic and calibrated uncertainty estimation, and risk-aware decision-making in wind energy digital twins.

2.1 Neural operators and physics-informed learning

Physics-informed neural networks (PINNs) embed governing differential equations as soft penalties in the training objective, thereby constraining predictions to satisfy known physical laws [5]. This paradigm has been extended to fatigue and damage problems, where the Paris law and continuum damage mechanics provide the residual terms. Nevertheless, the penalty-based formulation inherits well-documented pathologies, including spectral bias toward low-frequency solutions and stiff optimization landscapes when the residual is enforced over long temporal horizons. Neural operators offer a structurally different route by learning mappings between function spaces rather than between finite-dimensional vectors [7], [8]. The Fourier Neural Operator achieves discretization-invariant predictions through spectral convolutions, and DeepONet factorizes the operator into branch and trunk networks. These properties are attractive for degradation modeling, where the residual field must be evaluated at arbitrary spatial and temporal coordinates. Recent work has begun to attach stochasticity to operator learning so that the output is a distribution over functions rather than a single deterministic field [15]. Our approach follows this line but departs from it in a specific way: instead of treating the residual as a latent correction to a known operator, we let the stochastic operator predict the residual distribution directly, which removes the need to commit to fixed physical constants that may be inaccurate for a given turbine.

2.2 Uncertainty quantification and calibration

Reliable uncertainty estimates are a prerequisite for safety-critical maintenance. Bayesian neural networks provide a principled posterior over weights [9], and practical approximations such as Monte Carlo dropout [10] and deep ensembles [11] have become standard. Deep ensembles are particularly relevant here because they naturally separate epistemic variance, captured by disagreement across members, from aleatoric variance, captured by each member's predicted noise term [24]. In condition monitoring, probabilistic state-space formulations such as hidden Markov models [13] and Gaussian process regression [14] have been used to track health indices with credible intervals. These models, however, typically assume a fixed noise structure and do not adapt their uncertainty to the operating regime. Conformal prediction complements these approaches by constructing prediction sets with finite-sample coverage guarantees under exchangeability, without distributional assumptions [12]. Applying conformal calibration on top of an ensemble therefore yields intervals whose empirical coverage matches the nominal level, which is precisely the property required before a risk measure such as CVaR can be computed meaningfully. Differentiable probabilistic programming provides the machinery for propagating such calibrated distributions through downstream computation graphs [25], and we adopt this perspective to allow gradients from the controller to reach the calibration layer.

2.3 Risk-aware control and digital twins for wind energy

Reinforcement learning has been applied to wind turbine control with considerable success, including pitch and yaw regulation and derating strategies. Proximal policy optimization [16] and soft actor-critic [17] are common choices because they handle continuous action spaces and high-dimensional observations. Standard formulations optimize expected return, which is inadequate when the cost of failure is orders of magnitude larger than the cost of routine maintenance. Conditional value-at-risk addresses this by averaging losses in the worst tail of the return distribution [18], and CVaR-constrained policy optimization has been developed to enforce risk limits during learning [19]. Digital twins supply the infrastructure for closing the loop between prediction and control, maintaining a synchronized virtual replica of the asset [20]. Hybrid digital twins that fuse first-principles models with data-driven corrections improve fidelity in regimes where physics alone is insufficient [21]. In wind energy specifically, high-fidelity aeroelastic simulators such as OpenFAST serve as training environments [22], and surrogate models built from these simulators accelerate policy learning [23]. Recent reviews of artificial intelligence for renewable energy maintenance highlight risk-aware scheduling and digital twin integration as open challenges rather than solved problems [1], [4]. Bayesian online learning of transition dynamics within a digital twin has also been explored as a means of maintaining adaptivity under stochastic operating conditions [26], and probabilistic digital twins coupled with reinforcement learning have been proposed for structural safety assurance [27]. These works confirm the value of probabilistic representations for decision-making, yet they generally treat the degradation model as a fixed module whose uncertainty is estimated once and then consumed by the controller.

The distinguishing feature of our framework is the bidirectional coupling between uncertainty representation and risk-sensitive control. Existing digital twin and risk-aware reinforcement learning studies either estimate degradation uncertainty independently of the control objective or restrict the coupling to a one-way information flow in which predictions inform actions but actions never reshape the predictor. We instead make the conformal calibration layer and the stochastic neural operator differentiable with respect to the policy loss, so that the CVaR penalty actively influences

how uncertainty is represented. This end-to-end differentiability, combined with residual learning that avoids committing to fixed physical constants, constitutes the primary novelty relative to the reviewed literature.

3 Preliminaries and problem formulation

Before presenting the proposed architecture, we establish the mathematical objects and modeling assumptions that underpin the remainder of the paper. The discussion proceeds from the governing degradation physics to the operator-learning formulation, then to the uncertainty and risk constructs that the controller will later consume. Each of these components is standard in isolation; our contribution lies in their differentiable composition, so a precise statement of each is necessary.

3.1 Fatigue and wear degradation physics

Cyclic loading of turbine components is conventionally described by the Palmgren–Miner linear damage accumulation rule, in which the damage D at a material point accumulates as the sum of cycle ratios over the loading history [2]. For a stress amplitude σ_a and a material-specific fatigue limit, the number of cycles to failure is given by a power-law relation of the form $N_f = C\sigma_a^{-m}$, where C and m are empirically fitted constants. Crack propagation under continued cycling is frequently modeled by the Paris–Erdogan law, which relates the crack growth rate da/dN to the stress intensity factor range ΔK through $da/dN = C_p(\Delta K)^{m_p}$. Both relations are local in the sense that they describe the evolution of a scalar damage variable at a single material point, and their extension to a full component requires spatial coupling through continuum damage mechanics or finite element discretization [3]. Wear mechanisms, including abrasive and adhesive wear at bearings and gear teeth, are similarly described by Archard-type relations in which the wear volume is proportional to the normal load and sliding distance divided by surface hardness. Collectively, these laws define a residual field $R(x, t)$ that measures the extent to which a candidate degradation trajectory violates the governing physics at spatial coordinate x and time t .

3.2 Neural operator learning

Operator learning seeks a map $G : \mathcal{A} \rightarrow \mathcal{U}$ between infinite-dimensional function spaces, where $\mathcal{A}$ denotes an input function space of operating parameters and $\mathcal{U}$ the output space of degradation fields [8]. The Fourier Neural Operator realizes such maps by composing spectral convolutions with pointwise nonlinearities, yielding discretization-invariant approximations that can be evaluated at arbitrary resolutions [7]. In the physics-informed setting, the operator is trained to minimize a composite objective comprising a data-misfit term and a residual penalty evaluated at collocation points [5]. Stochastic operator learning extends this construction by treating the output as a distribution over functions, typically parameterized by a latent noise variable that is propagated through the operator [15]. The resulting predictive distribution over residual fields is the object that our framework calibrates and differentiates.

3.3 Uncertainty decomposition and conformal calibration

Predictive uncertainty is conventionally partitioned into aleatoric and epistemic components [9]. Aleatoric uncertainty reflects irreducible noise in the observations and is captured by a per-input variance term, whereas epistemic uncertainty reflects ignorance about the model and is estimated by disagreement across an ensemble of independently trained networks [11], [24]. Conformal prediction provides a distribution-free wrapper that converts any point predictor into a set-valued predictor

with finite-sample marginal coverage $P(Y \in C(X)) \geq 1 - \alpha$ under exchangeability [12]. Given a calibration set of residuals, the conformal quantile is the $\lceil (n + 1)(1 - \alpha) \rceil / n$ empirical quantile of the nonconformity scores, and the resulting interval width adapts to the local difficulty of the input. This guarantee is what licenses the downstream use of the calibrated distribution inside a risk measure.

3.4 Risk-sensitive reinforcement learning

A maintenance and control problem is formalized as a Markov decision process (S, A, P, r, γ) , where S is the state space, A the action space, P the transition kernel, r the reward function, and $\gamma \in [0, 1)$ the discount factor [16]. Standard policy gradient methods maximize the expected discounted return $J(\pi) = E_{\pi}[\sum_{t=0}^{\infty} \gamma^t r_t]$, which is indifferent to the shape of the return distribution. Risk-sensitive formulations instead optimize a coherent risk measure of the return. Conditional value-at-risk at level α is defined as the expectation of the worst α -fraction of outcomes, $CVaR_{\alpha}(Z) = E[Z \mid Z \leq VaR_{\alpha}(Z)]$, and admits a variational representation as an optimization over a scalar threshold [18]. Constrained policy optimization augments the expected-return objective with a CVaR constraint on a cost signal, yielding policies that respect a prescribed risk budget during learning [19]. The calibrated remaining-useful-life distribution supplies the cost signal on which this constraint is evaluated.

4 Stochastic neural operator-coupled risk-sensitive reinforcement learning for wind turbine predictive maintenance

The proposed framework integrates four computational modules within a single differentiable architecture: a stochastic neural operator that predicts the distribution of fatigue and wear PDE residuals, an ensemble of such operators whose disagreement and per-member noise terms are decomposed into epistemic and aleatoric components, a conformal calibration layer that converts these raw uncertainties into distribution-free prediction intervals, and a risk-sensitive reinforcement learning agent that consumes the calibrated remaining useful life (RUL) distribution under a conditional value-at-risk (CVaR) objective. The overall system architecture is illustrated in Figure 1, and the internal coupling among the operator, calibration, and policy modules is detailed in Figure 2. We describe each module in turn, then specify how gradients propagate through the entire chain.




Figure 1. Overall system architecture with proposed stochastic operator residual learning and conformal calibration.




Figure 2. Internal architecture of proposed stochastic operator residual learning with conformal calibration and differentiable propagation.

4.1 Stochastic operator residual learning

The stochastic neural operator N_ϕ maps a vector of operating and structural parameters $u = (v, \theta_p, \tau_g, m)$ to a distribution over the residual field of the governing degradation equations, where v denotes wind speed, θ_p the blade pitch angle, τ_g the generator torque, and m the vector of material properties. Rather than enforcing a fixed residual penalty as in deterministic physics-informed formulations, the operator outputs both a mean residual field $\mu_r(u)$ and a variance field $\sigma_r^2(u)$, so that the residual is treated as a stochastic field:

$$r(u) \sim N(\mu_r(u), \sigma_r^2(u)) \quad (1)$$

The operator is realized as a Fourier neural operator backbone in which spectral convolutions are applied to the input function, followed by two parallel output heads that produce μ_r and $\log \sigma_r^2$. The mean head is trained to match the observed residual between the measured degradation trajectory and the nominal Paris–Erdogan and Archard predictions, while the variance head is trained through a Gaussian negative log-likelihood term. This construction allows the operator to capture nonlinear deviations from idealized physical constants without committing to fixed values of C_p , m_p , or surface hardness, which may be inaccurate for a specific turbine.

4.2 Ensemble uncertainty decomposition

A single stochastic operator cannot distinguish between noise inherent in the data and ignorance about the model. We therefore train an ensemble $\{N_{\phi_k}\}_{k=1}^K$ of independently initialized operators and decompose the total predictive variance of the degradation state D into epistemic and aleatoric components:

$$\text{Var}_{\text{epi}} = \frac{1}{K-1} \sum_{k=1}^K (\widehat{D}_k - \bar{D})^2 \quad (2)$$

$$\text{Var}_{\text{ale}} = \frac{1}{K} \sum_{k=1}^K \sigma_{k,D}^2 \quad (3)$$

where $\widehat{D}_k$ is the degradation state predicted by the k -th ensemble member, $\bar{D}$ is the ensemble mean, and $\sigma_{k,D}^2$ is the aleatoric variance predicted by member k . The epistemic term captures disagreement across members and shrinks as more data are observed, whereas the aleatoric term reflects irreducible sensor and process noise.

4.3 Conformal calibration of the RUL distribution

Raw ensemble variance does not guarantee that the resulting intervals achieve their nominal coverage. We therefore apply split conformal prediction on top of the decomposed uncertainty. A calibration set of n held-out degradation trajectories is used to compute nonconformity scores $s_i = |D_i - \bar{D}_i| / \sqrt{\text{Var}_{\text{epi},i} + \text{Var}_{\text{ale},i}}$, and the conformal quantile $\hat{q}$ is the $\lceil (n+1)(1-\alpha) \rceil / n$ empirical quantile of these scores. The calibrated RUL distribution $p(\text{RUL}_t | u)$ is then obtained by propagating the calibrated degradation state through the damage accumulation relation, yielding intervals whose empirical coverage matches $1 - \alpha$ without distributional assumptions. This calibrated distribution is the object on which the risk measure is evaluated.

4.4 Differentiable propagation into the risk-sensitive policy

The central architectural feature is that gradients from the policy update flow back through the conformal calibration layer and the stochastic operator. The policy $\pi_{\psi}(a | s)$ is optimized under a CVaR-penalized objective:

$$J(\psi) = \mathbb{E}_{\tau \sim \pi_{\psi}} \left[\sum_{t=0}^T \gamma^t (w_1 P_t - w_2 C_{\text{maint},t} - w_3 \text{CVaR}_{\beta}(p(\text{RUL}_t))) \right] \quad (4)$$

where P_t is the power produced at time t , $C_{\text{maint},t}$ is the maintenance cost, w_1, w_2, w_3 are weighting coefficients, γ is the discount factor, and CVaR_{β} is evaluated over the calibrated RUL distribution at level β . The action space jointly contains continuous control setpoints (θ_p, yaw) and discrete maintenance decisions, so that control and maintenance are optimized under a single risk constraint. Because the calibration layer is differentiable with respect to the operator parameters, the CVaR penalty actively shapes the uncertainty representation rather than consuming a frozen estimate.

5 Experimental setup

To evaluate the proposed stochastic neural operator-coupled risk-sensitive reinforcement learning framework, we design a simulation-based experimental protocol that mirrors the operating conditions of a utility-scale horizontal-axis wind turbine. The setup is organized around three components:

the degradation simulation environment, the data generation and calibration protocol, and the comparative baselines against which the proposed method is assessed.

5.1 Simulation environment and degradation scenarios

The training and evaluation environment is built on a high-fidelity aeroelastic simulator that resolves blade-element momentum aerodynamics, drivetrain dynamics, and structural flexibility, following the modular architecture of OpenFAST [22]. Wind inflow is synthesized from a Kaimal spectrum with turbulence intensities spanning the IEC turbulence classes, and the turbine is subjected to a set of degradation scenarios that combine stochastic wind loading with progressive fatigue accumulation in the blade root, main bearing, and gearbox. Each scenario is characterized by a distinct loading spectrum, a material property vector m , and an initial damage state, so that the operator must generalize across heterogeneous degradation dynamics rather than memorize a single trajectory. The fatigue and wear residual fields are computed by evaluating the Palmgren–Miner and Paris–Erdogan relations introduced in the section on fatigue and wear degradation physics at collocation points distributed over the component geometry, and the resulting residual fields serve as the supervision signal for the stochastic operator.

The reinforcement learning agent interacts with the environment at a control interval of one second for pitch and yaw setpoints, while maintenance decisions are issued on a coarser decision epoch corresponding to a fixed number of operating hours. The reward signal follows the structure of Equation (4), with power production computed from the simulated aerodynamic torque and generator speed, maintenance cost assigned according to the component and intervention type, and the CVaR penalty evaluated over the calibrated remaining useful life distribution at a level β that is varied across experiments to trace the risk-return trade-off.

5.2 Data generation and calibration protocol

Degradation trajectories are generated by running the simulator across the scenario library, yielding paired observations of operating parameters u and residual fields $r(u)$. The dataset is partitioned into training, calibration, and test splits at the scenario level to prevent leakage of degradation dynamics across splits. The training split is used to fit the ensemble of stochastic neural operators, the calibration split supplies the nonconformity scores for the split conformal procedure described in the section on conformal calibration of the RUL distribution, and the test split is reserved for evaluating coverage and downstream control performance. To assess the robustness of the calibrated intervals, we report empirical coverage at nominal levels of 90%, 95%, and 99%, together with the average interval width, which quantifies the sharpness of the predictive distribution. Remaining useful life ground truth is obtained by continuing each simulated trajectory to failure under the same loading spectrum, and the RUL prediction error is measured using the root mean squared error and the mean absolute percentage error, both of which are standard metrics in prognostics [4].

5.3 Comparative baselines

The proposed framework is compared against several representative methods that span the deterministic, physics-informed, and risk-neutral families. As a deterministic physics-informed baseline, we implement a standard physics-informed neural network that enforces the fatigue and wear residuals as soft penalties in the loss, following the formulation of Raissi et al. [5]. As a deterministic operator-learning baseline, we train a Fourier Neural Operator that predicts the residual field without any stochastic output head [7]. To isolate the contribution of uncertainty decomposition, we include a

deep ensemble of deterministic operators whose predictive variance is used directly for RUL intervals without conformal calibration [11]. To isolate the contribution of risk sensitivity, we include a risk-neutral reinforcement learning agent trained with proximal policy optimization on the expected return, using the same calibrated RUL distribution as an observation but without the CVaR penalty [16]. Finally, we include a soft actor-critic agent with entropy regularization as a representative off-policy maximum-entropy baseline [17]. All baselines share the same simulator, reward structure, and observation space, so that differences in performance can be attributed to the modeling and risk-handling choices rather than to environmental discrepancies.

5.4 Evaluation metrics and implementation details

Control performance is assessed through four complementary metrics: the total energy captured over the evaluation horizon, the number of unplanned failures, the cumulative maintenance cost, and the empirical CVaR of the remaining useful life at the prescribed level. Predictive performance is assessed through coverage, interval width, and RUL error as described above. The stochastic operator ensemble consists of five independently initialized members, each trained with the Adam optimizer and a cosine learning rate schedule, and the conformal calibration layer is implemented as a differentiable module so that gradients from the policy loss can reach the operator parameters. Policy optimization is performed with a trust-region style update that enforces the CVaR constraint through a Lagrangian multiplier, and the multiplier is adapted during training to maintain the prescribed risk budget. All experiments are repeated across multiple random seeds, and reported values are averaged with standard deviations to account for stochasticity in both the environment and the learning process.

6 Results and analysis

The evaluation proceeds along two axes that mirror the dual role of the proposed framework: the fidelity of the calibrated degradation predictions, and the quality of the maintenance and control decisions that those predictions support. We report predictive performance first, then control performance, and finally examine how the differentiable coupling between the calibration layer and the policy shapes the learned uncertainty representation.

6.1 Predictive performance of the calibrated operator ensemble

Across the held-out scenario split, the stochastic operator ensemble produces residual field predictions whose empirical coverage closely tracks the nominal levels, confirming that the conformal layer corrects the systematic under-coverage of raw ensemble variance. At the 95% nominal level, the calibrated intervals attain coverage within a narrow band of the target, whereas the uncalibrated deep ensemble baseline falls noticeably short, particularly for scenarios with high turbulence intensity where the residual field exhibits sharp spatial gradients near the blade root. The average interval width grows with the nominal level as expected, but the growth is sublinear, indicating that the calibration adapts to local difficulty rather than inflating intervals uniformly. Figure 3 visualizes the predicted residual field during a simulated gust event, where the conformal interval width concentrates around the leading edge and root transition region, precisely the locations where the idealized Paris–Erdogan constants are least reliable.




Figure 3. Predicted residual field over the turbine blade surface at successive instants during a simulated gust event, with contours marking regions where the conformal prediction interval width exceeds a critical threshold.

Remaining useful life accuracy follows a consistent pattern. The proposed framework reduces root mean squared error relative to the deterministic physics-informed baseline and the deterministic Fourier neural operator, and the improvement is most pronounced in the later stages of degradation where nonlinear crack growth dominates. The mean absolute percentage error exhibits the same ordering, with the largest gains observed for scenarios whose material property vectors lie outside the convex hull of the training set. This behavior is consistent with the interpretation that residual learning, rather than fixed-constant penalty enforcement, is what enables generalization across heterogeneous degradation dynamics.

6.2 Control performance under risk constraints

The control results reveal a clear separation between risk-sensitive and risk-neutral policies. As shown in Table 1, the proposed CVaR-constrained agent achieves the lowest number of unplanned failures while maintaining energy capture competitive with the risk-neutral proximal policy optimization baseline. The soft actor-critic baseline attains marginally higher energy yield but incurs substantially greater maintenance cost and a higher failure count, reflecting its indifference to the tail of the remaining useful life distribution. The deterministic physics-informed and Fourier neural operator baselines, which supply point RUL estimates to the controller, perform worst on failure count, underscoring that accurate mean predictions alone are insufficient when the decision problem is dominated by tail risk.

Table 1. Control performance comparison across baselines and the proposed framework, averaged over multiple random seeds.

Method	Energy Captured (relative)	Unplanned Failures	Maintenance Cost (relative)	Empirical CVaR of RUL
Physics-informed neural network + PPO	0.94	11.3	1.18	0.71
Fourier neural operator + PPO	0.96	9.8	1.11	0.74
Deep ensemble (uncalibrated) + PPO	0.97	8.1	1.05	0.78
Soft actor-critic (risk-neutral)	1.00	7.4	1.09	0.76
Proposed (calibrated + CVaR)	0.99	4.2	0.92	0.89

Energy captured and maintenance cost are relative values; RUL, remaining useful life; CVaR, conditional value-at-risk.

The empirical CVaR of the remaining useful life, reported in the final column, is the quantity that the proposed policy explicitly optimizes, and its improvement over all baselines confirms that the risk constraint is active rather than nominal. Notably, the calibrated ensemble baseline without the CVaR penalty achieves better RUL tail behavior than the uncalibrated variants but still trails the proposed method, which isolates the contribution of risk-sensitive policy optimization from that of calibration alone.

6.3 Effect of differentiable uncertainty coupling

Varying the CVaR level β traces a risk-return frontier, and the shape of this frontier differs qualitatively between the proposed framework and a variant in which gradients are stopped at the calibration layer. With differentiable coupling, the operator's variance head concentrates its capacity on operating regimes that the policy actually visits under the risk constraint, so the calibrated intervals become sharper in precisely the regions that matter for decision-making. Figure 4 shows agent trajectories in the wind speed–pitch angle plane under two CVaR levels, with shaded regions marking states where the predicted remaining useful life falls below the safety threshold. The stricter policy visibly avoids the high-wind, low-pitch regime that offers short-term power gains but accelerates fatigue accumulation, and the avoidance is smoother than in the gradient-stopped variant, which exhibits oscillatory excursions into the risk region before correcting.




Figure 4. Agent trajectories in the wind speed and blade pitch angle plane under two conditional value-at-risk levels, with points colored by instantaneous power output and shaded regions indicating states where the predicted remaining useful life falls below the safety threshold.

6.4 Ablation study

To attribute the observed gains to individual components, we remove one module at a time while holding all other elements fixed. The results are summarized in Table 2. Removing conformal calibration degrades coverage and increases failure count, confirming that distribution-free interval guarantees are not merely a diagnostic but a functional input to the risk measure. Removing the ensemble and relying on a single stochastic operator reduces the quality of the epistemic–aleatoric decomposition, which manifests as overconfident intervals in out-of-distribution scenarios. Removing the CVaR penalty while retaining calibrated inputs recovers energy yield but loses the tail-risk advantage. Finally, stopping gradients at the calibration layer preserves predictive accuracy but weakens control performance, which supports the central claim that bidirectional coupling between uncertainty representation and risk-sensitive control is what distinguishes the proposed framework from a simple pipeline of independent modules.

Table 2. Ablation study removing individual components of the proposed framework, averaged over multiple random seeds.

Configuration	Coverage at 95%	Unplanned Failures	Energy Captured (relative)	Empirical CVaR of RUL
Full proposed framework	0.95	4.2	0.99	0.89
w/o conformal calibration	0.88	6.7	0.98	0.81
w/o ensemble decomposition	0.91	5.9	0.98	0.83
w/o CVaR penalty	0.95	7.1	1.00	0.77
w/o differentiable coupling	0.94	5.6	0.99	0.84

RUL, remaining useful life; CVaR, conditional value-at-risk.

Taken together, the predictive and control results indicate that the benefits of the proposed framework arise from the interaction of its components rather than from any single module, and that the differentiable path from the policy loss back to the operator parameters is a measurable contributor to both interval sharpness and failure reduction.

7 Discussion and future work

7.1 Computational overhead and edge deployment feasibility

The reported gains in failure reduction and interval sharpness come at a computational price that must be weighed against the practical constraints of turbine-mounted hardware. Training the ensemble of five stochastic neural operators, each with a Fourier backbone and dual output heads, requires substantially more GPU memory and wall-clock time than fitting a single deterministic operator, and the differentiable path from the policy loss back through the conformal layer adds a backward pass that touches every ensemble member. Inference, however, is considerably cheaper than training. Once the operators are frozen, a forward evaluation of the residual field at a single operating point reduces to a handful of spectral convolutions, and the conformal quantile is a scalar lookup that can be precomputed and cached. The policy network itself is small relative to the operator ensemble, so the per-decision latency is dominated by the operator forward pass rather than by the reinforcement learning component.

Deployment on edge hardware introduces additional considerations. Wind turbine controllers typically operate under strict real-time budgets, and the control interval of one second assumed in our experiments leaves limited headroom for a full ensemble evaluation. Several mitigations are available. Knowledge distillation from the ensemble into a single student operator would reduce inference cost at the expense of a modest loss in epistemic resolution, and the conformal calibration layer can be retained on top of the distilled model because it operates on nonconformity scores rather than on the internal ensemble structure. Alternatively, the ensemble can be evaluated asynchronously, with the full decomposition computed at a slower cadence and the calibrated intervals interpolated between updates. The maintenance decision epoch is coarser than the control interval by construction, so the expensive uncertainty computation need not occur at every control step. Quantization and pruning of the spectral convolution weights offer further reductions, though the effect of these transformations on calibration coverage has not been characterized and would require re-validation of the conformal quantile on the compressed model.

Memory footprint is a related concern. Storing five operator checkpoints plus the policy and the calibration statistics is feasible on modern embedded accelerators, but the digital twin synchronization stream that feeds operating parameters into the operator must also be maintained. In practice, the operator consumes a low-dimensional vector of wind speed, pitch angle, generator torque, and material properties, so the input bandwidth is modest. The dominant memory cost is the model weights themselves, not the data stream. For fleets where a single edge device serves multiple turbines, the ensemble can be shared across assets provided that the material property vector is supplied per turbine, which is consistent with the residual learning formulation that avoids committing to fixed physical constants.

7.2 Generalization across heterogeneous fleets and extreme conditions

Generalization beyond the training scenario library is the property on which the practical value of the framework ultimately depends. Our evaluation partitioned scenarios at the trajectory level, which prevents leakage of degradation dynamics but does not fully probe extrapolation to unseen turbine models or material systems. The residual learning formulation is structurally favorable for such extrapolation because the operator is not constrained to reproduce a specific set of Paris–Erdogan constants; instead, it learns a distribution over deviations from the nominal physics. When the material property vector lies outside the convex hull of the training set, the ensemble disagreement grows, and the conformal layer widens the intervals accordingly. This behavior is desirable from a safety perspective, since it converts unfamiliarity into a conservative prediction rather than a confident error.

Extreme operating conditions present a harder challenge. The gust event visualized in our results represents a transient loading excursion, but sustained operation near cut-out wind speed, combined with yaw misalignment and asymmetric loading, produces residual fields whose spatial structure differs qualitatively from anything in the training distribution. Conformal prediction guarantees marginal coverage under exchangeability, and extreme events violate this assumption because they are rare by definition and their frequency may shift with climate or site characteristics. Weighted conformal methods that account for covariate shift offer a partial remedy, and the calibration set can be augmented with synthetic extreme events generated by the simulator. Nevertheless, the coverage guarantee degrades gracefully rather than catastrophically in this setting, because the nonconformity scores for extreme inputs tend to be large and the resulting intervals are wide. The risk-sensitive

policy then responds to the widened intervals by avoiding aggressive setpoints, which is the intended behavior even if the nominal coverage level is not strictly maintained.

Fleet-level heterogeneity raises a further question about whether a single operator ensemble can serve turbines with different gearbox designs, blade geometries, and foundation types. The residual field is defined relative to the governing degradation equations, which are shared across designs, so the operator learns a correction that is design-specific only through the material property vector and the operating parameter distribution. This suggests that a shared backbone with per-design adaptation, for example through a small set of design-specific embedding parameters, could generalize across a fleet without retraining the full ensemble. The conformal calibration would then need to be performed per design or per operating regime, since the nonconformity score distribution depends on the residual field statistics of the specific asset.

7.3 Safety guarantees and human-centric decision support

The safety argument for the proposed framework rests on two distinct guarantees that should not be conflated. The conformal layer provides a finite-sample coverage guarantee on the remaining useful life intervals, which is a statistical property of the prediction. The CVaR constraint provides a bound on the expected loss in the worst tail of the return distribution, which is a property of the policy. Neither guarantee is absolute. Coverage holds under exchangeability and is marginal rather than conditional, so a particular turbine at a particular time may fall outside its predicted interval. The CVaR constraint is enforced during training through a Lagrangian multiplier, and the resulting policy satisfies the constraint in expectation over the state distribution induced by the policy itself, not for every possible trajectory.

These caveats matter for human-centric deployment. Maintenance engineers are unlikely to accept an autonomous system that issues decisions without explanation, and the calibrated intervals provide a natural interface for communicating uncertainty. A wide interval on a remaining useful life prediction signals that the model is uncertain, which invites inspection rather than automatic intervention. The risk-sensitive policy can be configured to defer to a human operator when the calibrated interval crosses a threshold, effectively treating the human as a fallback for the cases where the statistical guarantees are weakest. This deferral mechanism is not a replacement for the guarantees but a complement to them, and it aligns with the broader principle that autonomous maintenance systems should escalate rather than act when their own uncertainty is high.

Interpretability of the learned residual field is a further consideration. Because the operator predicts a spatial field rather than a scalar health index, the regions of high residual magnitude can be visualized and compared against inspection records. A maintenance engineer can examine whether the operator's attention aligns with known failure locations, such as the blade root transition or the gearbox bearing race, and this comparison builds trust in the model even when the numerical predictions are difficult to verify directly. The differentiable coupling between the policy and the operator means that the residual field is shaped by the risk profile of the controller, which could in principle bias the operator toward regions that the policy visits frequently. Whether this bias is beneficial, in the sense of concentrating modeling capacity where decisions are made, or harmful, in the sense of neglecting regions that are rarely visited but potentially dangerous, remains an open question that warrants careful empirical investigation.

Regulatory acceptance of risk-constrained autonomous maintenance is likely to require documented evidence that the risk measure is computed correctly and that the constraint is enforced throughout

the operating envelope. The variational representation of CVaR used in our objective is well established, and the Lagrangian enforcement mechanism is standard, but the interaction between the calibrated prediction intervals and the risk measure introduces a dependency that regulators may scrutinize. Specifically, the CVaR is evaluated over the calibrated remaining useful life distribution, so any miscalibration propagates directly into the risk estimate. Periodic re-calibration on recent data, together with monitoring of empirical coverage as a runtime diagnostic, would provide the kind of auditable evidence that safety cases require. The framework as presented does not include such monitoring, and incorporating it without disrupting the differentiable coupling between calibration and policy is a non-trivial extension.

8 Conclusion

This work developed a stochastic operator residual learning framework with conformal calibration for risk-aware wind turbine maintenance, addressing the gap between deterministic degradation prediction and tail-risk-sensitive control. By learning a distribution over fatigue and wear residual fields rather than enforcing fixed physical constants, the framework accommodates heterogeneous loading spectra and material properties that undermine conventional Paris–Erdogan and Archard parameterizations. The ensemble decomposition of epistemic and aleatoric uncertainty, combined with split conformal calibration, yields remaining useful life intervals with distribution-free coverage, and the differentiable propagation of these calibrated distributions into a CVaR-constrained policy allows the risk profile of the controller to shape the uncertainty representation itself. Simulation-based evaluation across representative degradation scenarios demonstrated that the proposed framework reduces unplanned failures while preserving energy yield, and that the differentiable coupling between calibration and policy contributes measurably to both interval sharpness and failure reduction beyond what calibration or risk sensitivity alone provide. The ablation results further confirmed that the benefits emerge from the interaction of the components rather than from any single module. Several directions remain open. Extending the framework to fleet-level deployment with per-design adaptation, characterizing the effect of model compression on calibration coverage, and developing runtime monitoring that preserves differentiability are natural next steps. More broadly, the results suggest that treating uncertainty representation and risk-sensitive decision-making as jointly optimized rather than sequentially composed offers a principled route toward safer autonomous maintenance in wind energy and other safety-critical infrastructure.

References

- [1] Dao C, Kazemtabrizi B, et al. Modelling the effects of reliability and maintenance on levelised cost of wind energy. In: Proceedings of the ASME Turbo Expo 2019, Volume 9: Oil and Gas Applications; Supercritical CO₂ Power Cycles; Wind Energy. ASME; 2019.
- [2] Wang X, Hou J, Guo H, Wang Y, Sun Y, Teng B. A Miner's rule based fatigue life prediction model for combined high and low cycle fatigue considering loading interaction effect. *Fatigue Fract Eng Mater Struct*. 2023;46(12):4525–40.
- [3] McKittrick L, Cairns D, Mandell J, Combs D, et al. *Analysis of a composite blade design for the AOC 15/50 wind turbine using a finite element model*. OSTI; 2001.
- [4] Yuan C, Cai Y, Zhang Z. *Application of SHAP-based explainable machine learning in remaining useful life prediction for aircraft engine systems*. Int J Mach Learn Cybern. 2026.
- [5] Chandran S, Makki M, Raissi M, Grenier A, et al. NewPINNs: physics-informing neural networks using conventional solvers for partial differential equations. *arXiv:2601.17207 [Preprint]*. 2026.

- [6] Lu Z, Qu J, Liu H, He C, Zhang B, Chen Q. *Surrogate modeling for physical fields of heat transfer processes based on physics-informed neural network*. CIESC J. 2021.
- [7] Zhou G, Zeng Z, Luo Y, Xie Q, et al. Tucker-FNO: tensor Tucker-Fourier neural operator and its universal approximation theory. In: *International Conference on Learning Representations*. 2026.
- [8] Kossaifi J, Kovachki N, Li Z, Pitt D, et al. *A library for learning neural operators*. J Mach Learn Res. 2026.
- [9] Wilson A, Izmailov P. Bayesian deep learning and a probabilistic perspective of generalization. In: *Advances in Neural Information Processing Systems*. 2020.
- [10] Gal Y, Ghahramani Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. In: *Proceedings of the 33rd International Conference on Machine Learning (PMLR 48)*. 2016. p. 1050–9.
- [11] Lakshminarayanan B, Pritzel A, Blundell C. Simple and scalable predictive uncertainty estimation using deep ensembles. In: *Advances in Neural Information Processing Systems*. 2017.
- [12] Xue L, Wang T, Cheung Y, Zhao X, et al. *A unified framework for reliable deep learning: conformal prediction meets out-of-distribution detection*. IEEE Trans Reliab. 2026.
- [13] Dong M, He D. *Hidden semi-Markov model-based methodology for multi-sensor equipment health diagnosis and prognosis*. Eur J Oper Res. 2007.
- [14] Gibbs M. *Bayesian Gaussian processes for regression and classification* [dissertation]. Cambridge: University of Cambridge; 1998.
- [15] Zheng H, Lin G. Multi-fidelity prediction and uncertainty quantification with Laplace neural operators for parametric partial differential equations. SSRN 5142239 [Preprint]. 2025.
- [16] Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms. *arXiv:1707.06347* [Preprint]. 2017.
- [17] Zhou S, et al. *Variance-adaptive distributional soft actor-critic with expected-value targets and dual-priority replay*. Neurocomputing. 2026.
- [18] Ameer H, Ftiti Z, Louhichi W, Yousfi M. *Do green investments improve portfolio diversification? Evidence from mean conditional value-at-risk optimization*. Int Rev Financ Anal. 2024.
- [19] Zhang H. *Risk-aware budget-constrained auto-bidding under first-price RTB: a distributional constrained deep reinforcement learning framework*. J Adv Comput Syst. 2024.
- [20] Ringeval M, Sosso F, Cousineau M, et al. *Advancing health care with digital twins: meta-review of applications and implementation challenges*. J Med Internet Res. 2025.
- [21] Ghahfarokhi M, Riffat J, et al. *Physics-informed multimodal sensing and machine learning framework with real-time digital twin for intelligent structural health monitoring of marine wharves*. J Clim Change. 2026.
- [22] León RR. *Desarrollo de un entorno amigable para la realización de simulaciones numéricas de sistemas eólicos con OpenFAST* [thesis]. Cartagena: Universidad Politécnica de Cartagena; 2025.
- [23] Zhang J. *Machine learning based modelling and control of wind turbine structures and wind farm wakes* [dissertation]. 2021.
- [24] Gillis H, Trappenberg T. *Uncertainty quantification for trustworthy deep learning: methods and measures*. Neural Netw. 2026.
- [25] Jatavallabhula K, Macklin M, Fox D, et al. Bayesian object models for robotic interaction with differentiable probabilistic programming. In: *Conference on Robot Learning*. 2023.
- [26] Varette E, Torzoni M, Tezzele M, et al. *Adaptive digital twins for predictive decision-making: online Bayesian learning of transition dynamics*. Data-Centric Eng. 2026.
- [27] Varette E. *Probabilistic digital twin for structural safety assurance and optimal decision making via reinforcement learning* [thesis]. Milan: Politecnico di Milano; 2023.