



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Salinas-Camus, M., Patrick, M. A., Pascoe, J. A., & Eleftheroglou, N. (2026). *Bayesian Post-Repair Prognostics for Reliable RUL Prediction*. 1-8. Paper presented at Proceedings of the 9th European Conference of the Prognostics and Health Management Society 2026, Oslo, Norway. <https://www.papers.phmsociety.org/index.php/phme/issue/view/78>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Bayesian Post-Repair Prognostics for Reliable RUL Prediction

Mariana Salinas-Camus¹, Mary Patrick², John-Alan Pascoe³ and Nick Eleftheroglou⁴

^{1,4} *Intelligent System Prognostics for Operations and Maintenance Group, Department of Aerospace Structures and Materials, Faculty of Aerospace Engineering, Delft University of Technology, Delft, 2629HS, Netherlands.*

m.salinascamus@tudelft.nl

n.eleftheroglou@tudelft.nl

^{2,3} *Department of Aerospace Structures and Materials, Faculty of Aerospace Engineering, Delft University of Technology, Delft, 2629HS, Netherlands.*

m.a.patrick@tudelft.nl

j.a.pascoe@tudelft.nl

ABSTRACT

Prognostics aims to predict the Remaining Useful Life (RUL) of engineering assets and is essential for effective Predictive Maintenance (PdM). Unlike preventive maintenance, PdM offers substantial cost benefits by scheduling maintenance only when needed. However, most existing prognostic models assume that repairs return an asset to an “as-good-as-new” condition. In practice, repairs are often imperfect, as they only partly restore the asset and may change its subsequent degradation behavior. This mismatch represents a major limitation of current prognostic approaches, as poor prognostic performance can lead to unnecessary maintenance actions or unexpected failure.

This paper proposes a fully Bayesian prognostic model, named the Sequential Bayesian Semi-Markov Framework (SBSM), that explicitly accounts for imperfect repair while being trained exclusively on data from non-repaired assets. The framework combines a Hidden Semi-Markov Model (HSMM) to represent the degradation model with a particle filter for the predictive step. Repair actions are incorporated as prior distributions that represent repair effectiveness, enabling repair uncertainty and prognostic uncertainty to be treated in a unified manner. This formulation allows post-repair degradation trajectories to differ from pre-repair behavior without retraining the model.

The approach is experimentally validated using an in-house dataset of aluminum open-hole specimens subjected to constant-amplitude fatigue loading, repaired via cold spray deposition, and tested to failure. Baseline (non-repaired) specimens are used for training, while repaired specimens are used exclu-

sively for testing. The proposed method is compared against a Convolutional Neural Network (CNN) baseline. Results show that the SBSM achieves lower prediction error and improved probabilistic calibration, particularly under significant distributional shifts induced by more effective repairs. The framework demonstrates robust post-repair RUL prediction and well-calibrated uncertainty estimates, highlighting its potential for real-world predictive maintenance applications involving imperfect repair.

1. INTRODUCTION

Prognostics is a key component of predictive maintenance, as it provides Remaining Useful Life (RUL) predictions. Unlike reactive maintenance (after failure) or preventive maintenance (at fixed intervals), predictive maintenance schedules interventions according to the actual degradation state of the asset. This condition-based strategy improves reliability, reduces downtime, and optimizes life cycle cost (Moleda et al., 2023).

Maintenance actions can be classified as perfect, minimal, worse, or imperfect repair (Carlo & Arleo, 2017). A perfect repair restores the asset to an “as good as new” condition. A minimal repair restores it to an “as bad as old” state, leaving degradation unchanged. Worse repair degrades the asset further than its pre-repair state. Imperfect repair lies between perfect and minimal repair: the asset is partially restored but not to its original condition. In practice, most repairs are imperfect. If this is not properly modeled, RUL predictions may become overly optimistic, leading to unexpected failures (Do et al., 2015). Therefore, prognostic frameworks must explicitly account for imperfect repair and its uncertainty.

In prognostics, data-driven models are among the most widely used approaches, as they rely on historical data for training

Mariana Salinas-Camus et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

and do not require explicit physical knowledge of the asset. This is particularly advantageous for complex engineering assets, where obtaining accurate and complete physics-based models can be challenging and often impractical (Salinas-Camus et al., 2025). However, only a limited number of data-driven models incorporate imperfect repairs. Existing approaches either rely on computationally expensive deep learning architectures (Skordilis & Moghaddass, 2020), exhibit early-life prediction bias (Ma et al., 2023), or treat repair uncertainty independently from prognostic uncertainty without allowing degradation dynamics to change after repair (Komninos et al., 2025). As a result, current methods may fail to capture the complex interactions between repair actions and future degradation behavior.

To address these limitations, this paper proposes a fully Bayesian data-driven prognostic framework, called the Sequential Bayesian Semi-Markov Framework (SBSM) that explicitly incorporates imperfect repair and its uncertainty into RUL prediction. The proposed framework models the degradation process using a Hidden Semi-Markov Model (HSMM) (Yu, 2010), as it provides an effective representation of degradation based on condition monitoring (CM) data, enabling accurate RUL prediction while quantifying uncertainty (Dong & He, 2007; Kontogiannis et al., 2025; Wang et al., 2014). To predict RUL, the proposed SBSM uses sequential Bayesian inference to update the asset state as new observations become available. Repair actions are integrated into the prognostic process as probabilistic interventions, allowing the degradation trajectory to change after maintenance while accounting for uncertainty in repair effectiveness. This enables joint propagation of prognostic and repair uncertainty and allows the framework to operate without requiring repaired data during training.

The framework is validated through an in-house experimental study of aluminum specimens subjected to fatigue loading and repaired using cold spray, an additive manufacturing technique in which metallic powder particles are accelerated to supersonic velocities to form dense, adherent coatings at relatively low temperatures. Cold spray has applications in aerospace, automotive, and marine industries (Yandouzi et al., 2014; Widener et al., 2016). Baseline (non-repaired) specimens are used for training, while repaired specimens are used exclusively for testing. The proposed approach is further compared with a Convolutional Neural Network (CNN) as a deep learning baseline.

The main contributions of this work are summarized as follows:

1. A fully Bayesian prognostic framework that integrates HSMM-based degradation modeling with particle filtering to provide probabilistic RUL estimation.
2. A principled mechanism to incorporate imperfect repair

within the inference process, enabling repair uncertainty and prognostic uncertainty to be treated in a unified manner.

3. Experimental validation on cold spray repaired aluminum fatigue specimens, including comparison with a CNN-based deep learning baseline.

The remainder of this paper is organized as follows. Section 2 describes the methodology, detailing the SBSM in both offline and online phases, as well as the baseline CNN architecture used for comparison. Section 3 presents the case study on aluminum open-hole specimens repaired with cold spray. Section 4 reports the results of the SBSM and CNN in terms of prediction accuracy and uncertainty calibration. Finally, Section 5 summarizes the conclusions.

2. METHODOLOGY

The Sequential Bayesian Semi-Markov Framework (SBSM) operates in two stages: an offline (train) phase and an online (test) phase, as presented in Figure 1. During the offline phase, historical CM data or a Health Indicator (HI) are used to train the parameters of an HSMM, including the state-dependent observation distributions and the Weibull state duration parameters. The HSMM was implemented using the HiMAP package (Kontogiannis et al., 2026). The HSMM outputs the learned parameter set, θ , which characterizes the stochastic degradation process of the asset. These estimated parameters are subsequently used for Bayesian inference.

In the online phase, real-time observations are assimilated using a Particle Filter (PF). The PF is employed for sequential Bayesian inference because it provides a computationally feasible method for approximating the posterior distribution of the asset states (Jouin et al., 2016). Initialized with the learned HSMM parameters θ , the PF recursively estimates the posterior distribution of the hidden degradation states as new observations become available. In addition, the PF enables the incorporation of repair actions by introducing a prior distribution that represents repair effectiveness. This prior information is used to update the state duration distributions in a Bayesian manner. This combination of HSMM-based modeling and particle filtering provides a fully data-driven and probabilistic framework for reliable prognostics.

2.1. Hidden Semi-Markov Degradation Model

The degradation process of the engineering asset is modeled as an HSMM. Let $X_t \in \{1, \dots, S\}$ denote the hidden degradation state at time t , and let Y_t represent the corresponding CM observation.

Unlike standard hidden Markov models, the duration of the state is explicitly modeled through state-dependent duration distributions. For each degradation state s , the sojourn time follows a Weibull distribution parameterized by a shape pa-




Figure 1. Sequential Bayesian Semi-Markov Framework (SBSM) overview.

parameter α_s and a scale parameter β_s .

2.2. State Duration and Transition Model

Let e denote the elapsed time spent in the current state s . The survival function associated with the Weibull duration model is given by

$$S_s(e) = \exp \left[- \left(\frac{e}{\beta_s} \right)^{\alpha_s} \right]. \quad (1)$$

The conditional probability of remaining in state s for one additional time step is therefore expressed as

$$P(X_{t+1} = s \mid X_t = s, e) = \frac{S_s(e+1)}{S_s(e)} \quad (2)$$

$$= \exp \left[\left(\frac{e}{\beta_s} \right)^{\alpha_s} - \left(\frac{e+1}{\beta_s} \right)^{\alpha_s} \right]. \quad (3)$$

Accordingly, the probability that the current state terminates at elapsed time e is given by

$$P(X_{t+1} = s+1 \mid X_t = s, e) = 1 - \frac{S_s(e+1)}{S_s(e)}. \quad (4)$$

Transitions are assumed to be sequential, reflecting the monotonic progression of degradation through the hidden states.

2.3. Observation Model

Conditional on the hidden degradation state, observations are assumed to be independent and normally distributed. The emission model is defined as

$$p(Y_t \mid X_t = s) = \mathcal{N}(Y_t \mid \mu_s, \sigma_s^2), \quad (5)$$

where μ_s and σ_s denote the state-dependent mean and standard deviation, respectively. This assumption reflects the HSMM degradation model, where each state generates observations according to its emission parameters.

2.4. Initial Degradation Condition

Prior knowledge regarding the initial degradation condition of the asset is incorporated through a categorical prior distribution over the hidden states,

$$P(X_1 = s) = \pi_s, \quad (6)$$

where π_s represents the prior probability that the asset starts in state s .

2.5. Bayesian Filtering and Posterior Approximation

The objective of inference is to recursively estimate the filtering distribution

$$p(X_t \mid Y_{1:t}), \quad (7)$$

which is analytically intractable due to the explicit modeling of state durations. To address this challenge, the particle filtering approach is employed to approximate the posterior distribution.

Each particle represents a realization of the hidden state and its associated elapsed time. The particle weights are proportional to the joint likelihood of the observation and the transition dynamics,

$$\omega_t^{(i)} \propto p(Y_t \mid X_t^{(i)}) p(X_t^{(i)} \mid X_{t-1}^{(i)}, e_{t-1}^{(i)}). \quad (8)$$

After normalization, the filtering distribution is approximated

as

$$p(X_t = s \mid Y_{1:t}) \approx \sum_{i=1}^N \omega_t^{(i)} \mathbb{I}(X_t^{(i)} = s), \quad (9)$$

where $\mathbb{I}(\cdot)$ denotes the indicator function.

2.6. Particle Degeneracy Control

To mitigate particle degeneracy, the effective sample size is monitored during inference and defined as

$$N_{\text{eff}} = \frac{1}{\sum_{i=1}^N (\omega_t^{(i)})^2}. \quad (10)$$

When N_{eff} falls below a predefined threshold, resampling is performed to maintain particle diversity and ensure numerical stability of the approximation.

2.7. Incorporation of Repair Actions

The SBSM framework enables the incorporation of repair actions through an explicit probabilistic transition mechanism within the particle filtering process. Repair events are treated as exogenous interventions that alter the future degradation dynamics of the asset while preserving uncertainty regarding their effectiveness.

A repair action is assumed to either restore the asset to a previous hidden degradation state or to leave the current state unchanged while partially resetting the elapsed time variable e_t . The emission model remains unaffected by repair actions, as the relationship between the hidden degradation state and the observations is assumed to be the same. In contrast, the state duration distributions are allowed to change in order to reflect the modified degradation rate induced by the repair.

State durations are modeled using Weibull distributions with shape parameter α_s and scale parameter β_s (see Subsection 2.2), learned during the offline training of the HSMM. The shape parameter primarily governs the variability of the sojourn time, whereas the scale parameter controls its expected magnitude and thus directly influences the RUL prediction.

It is assumed that repair effectiveness can be quantified either through a diagnostic technique (e.g., non-destructive testing) or through a repair effectiveness model. This assumption is reasonable and applicable in real-time settings, as repair assessment is commonly performed immediately after maintenance actions to evaluate structural integrity or expected performance recovery. Therefore, this information regarding the effectiveness of a repair is incorporated by introducing a stochastic scaling factor applied to the Weibull scale parameters. Specifically, upon the occurrence of a repair action, a latent repair effectiveness variable r is sampled from a prior

distribution,

$$r \sim \mathcal{N}(\mu_r, \sigma_r^2), \quad (11)$$

where $\mu_r > 1$ reflects the expected improvement in degradation resistance due to repair, and σ_r^2 encodes uncertainty in the expert assessment.

Conditioned on the sampled value of r , the Weibull scale parameter of each particle is updated according to

$$\beta_s^+ = r \beta_s^-, \quad (12)$$

where β_s^- and β_s^+ denote the scale parameters immediately before and after the repair event, respectively.

As a result, repair actions introduce controlled epistemic uncertainty into the particle ensemble, allowing post-repair degradation trajectories to evolve heterogeneously across particles. Subsequent observations progressively constrain this uncertainty through the standard weighting and resampling steps of the particle filter.

2.8. Remaining Useful Life prediction

The asset's RUL is predicted by propagating the particle ensemble, explicitly accounting for both aleatoric and epistemic uncertainties. For each particle, the current elapsed time in the degradation state is considered, and the remaining time in that state is sampled from the corresponding Weibull distribution conditioned on survival beyond the elapsed duration:

$$t_{\text{rem}} = \text{Weibull}(\alpha_s, \beta_s \mid t > e) - e, \quad (13)$$

where s is the current state and e is the elapsed time. Following this, the durations of subsequent states are sampled sequentially according to their respective Weibull distributions. This generates a particle level trajectory of the total remaining life:

$$\text{RUL}^{(i)} = t_{\text{rem}}^{(i)} + \sum_{s' > s} t_{s'}^{(i)}, \quad (14)$$

where $t_{s'}^{(i)}$ represents a sampled sojourn time for future state s' for particle i .

To more accurately represent the inherent stochasticity of the degradation process, multiple realizations of the remaining life are generated for each particle. Using only a single sample per particle would provide only one possible trajectory, potentially underestimating the variability due to the randomness in state durations. By generating multiple trajectories per particle, the model captures the aleatoric uncertainty in the asset, while the particle distribution itself reflects epistemic uncertainty arising from partial knowledge of the current state or the effects of repair actions.

2.9. Baseline model: Convolutional Neural Network

To evaluate the performance of the SBSM framework, a Convolutional Neural Network (CNN) is used as the baseline model, as CNNs have been widely adopted in prognostics with the rise of deep learning. The CNN is trained to map fixed-length temporal windows of CM data to the RUL. Its architecture is illustrated in Figure 2. The extracted feature representations are then mapped to a scalar output via a fully connected layer with a linear activation. Network parameters θ are optimized by minimizing the mean squared error (MSE) between predicted and true RUL values:

$$\mathcal{L}(\theta) = \frac{1}{N} \sum_{i=1}^N (y_i - f_{\theta}(\mathbf{X}_i))^2, \quad (15)$$

where N is the number of training windows and $f_{\theta}(\mathbf{X}_i)$ denotes the CNN prediction for input $\mathbf{X}_i$.




Figure 2. CNN architecture.

To capture predictive uncertainty, Monte Carlo (MC) Dropout is applied during both training and inference. Dropout layers remain active during prediction, producing stochastic forward passes that generate a set of RUL predictions (samples) for each input window. These samples create an empirical RUL distribution and represent an approximation of the epistemic uncertainty (Gal & Ghahramani, 2016).

Unlike the particle filter HSMM, the CNN does not explicitly model the asset's degradation process. Temporal dependencies are learned implicitly through convolutional feature extraction, while MC Dropout provides a probabilistic approximation of epistemic uncertainty.

3. CASE STUDY

Specimens were fabricated from Aluminium 6082-T6 (300 mm × 60 mm × 3 mm) and featured a 6 mm diameter hole.

A 0.5 mm long notch with a 0.125 mm radius was machined on each side of the hole using wire electric discharge machining (EDM), oriented transverse to the loading axis to promote localized crack initiation.

Fatigue testing was performed using load control in accordance with ASTM E466 (ASTM International, 2021). Testing was performed using a hydraulic machine with a 100 kN load cell under constant amplitude loading, with a peak load of 13 kN, a load ratio of 0.1, and a frequency of 5 Hz. After each series of 250 cycles, a 2-second hold period at the maximum load was included to capture images for Digital Image Correlation (DIC), enabling strain measurement as an indicator of damage progression.

Training specimens were tested to failure, whereas testing specimens were fatigued to 45,000 cycles before undergoing cold spray repair. Before deposition, the repair region was locally abraded with Scotch-Brite and cleaned with isopropyl alcohol. Spraying was conducted within four hours of surface preparation to limit oxide formation. Cold spraying was performed using nitrogen as the carrier gas at 0.2 MPa (20 bar) with a nozzle coil temperature of 600 °C. The nozzle was positioned at a 20 mm stand-off distance and traversed the coupons at 20 mm/s with a raster step size of 0.5 mm (center-to-center distance between adjacent passes). Two repair configurations were applied: thin repairs, consisting of a single deposited layer (~0.5 mm total thickness), and thick repairs, consisting of three layers (~1.5 mm total thickness).




Figure 3. Non-repaired specimen exhibiting crack growth (left), and repaired specimen (right) with cold spray patches applied over the crack region.

Specimens 1, 2, 3, 5, and 8 received a thick repair, while specimens 4, 6, 7, and 9 received a thin repair. After repair, all specimens were subjected to the same loading conditions as before the repair until failure, with repair thickness influencing total lifetimes (see Figure 4).

For preprocessing, the first 35,000 cycles of each degradation history were removed to exclude the initial flat ("healthy") stage. The strain measurements were then transformed into a cumulative feature and discretized to obtain an idealized




Figure 4. Strain data from training and testing sets.

(“perfect”) feature. Although this feature is not suitable for real-time monitoring, it serves as a benchmark for evaluating the prognostic model performance.

The training set consisted of 10 degradation histories from non-repaired specimens, with an average lifetime of 56.50 kcycles (91.50 kcycles total before pre-processing). The test set included 9 histories: 5 with thick repairs and 4 with thin repairs. Thick repairs had a longer average fatigue life of 80.60 kcycles (115.60 kcycles total before pre-processing), whereas thin repairs resulted in an average lifetime of 61.90 kcycles (96.93 kcycles total before pre-processing).

4. RESULTS

The SBSM framework requires training the HSMM, for which the Bayesian Information Criterion (BIC) was used to select the optimal number of hidden states. The learned HSMM parameters θ , representing the degradation model and trained exclusively on non-repaired specimens, were then utilized within the PF, which used 100 particles to obtain probabilistic RUL predictions. In this study, prior knowledge about repair effectiveness was introduced through a Gaussian distribution. For thin repairs, it was assumed that the lifetime improves by 5% relative to the baseline degradation model, leading to $\mu_r = 1.05$ with a standard deviation of 0.02. For thick repairs, a 30% improvement in lifetime was assumed, resulting in $\mu_r = 1.3$ with a standard deviation of 0.05. These priors were injected into the particle ensemble at the repair time, allowing both the expected improvement and its associated uncertainty to be propagated during prediction.

In contrast, the hyperparameters of the CNN baseline were tuned through grid search, and the network was trained using sliding windows of length 10.

Performance was evaluated using RMSE to quantify point prediction accuracy and CRPS, to jointly assess prediction accuracy and uncertainty calibration and sharpness. Since the repair altered the End of Life (EoL) and therefore the true RUL trajectory, both metrics were computed only from

the repair point onward, ensuring that the comparison reflects post-repair prognostic performance.

Table 1 summarizes the prognostic performance for each specimen, with specimens that had a thick repair indicated by an asterisk. Overall, the results demonstrate that the SBSM consistently outperforms the CNN across all cases. The mean RMSE decreases from 7.52 kcycles (CNN) to 3.30 kcycles (SBSM), while the mean CRPS decreases from 4.96 kcycles to 2.15 kcycles. The improvement in CRPS is particularly important, as it demonstrates that the SBSM not only improves point prediction accuracy but also provides better predictive distributions.

Table 1. RMSE and CRPS for CNN and SBSM across test set.

Specimen	RMSE [kcycles]		CRPS [kcycles]	
	CNN	SBSM	CNN	SBSM
1*	10.08	1.91	6.41	1.70
2*	11.44	2.88	7.42	2.09
3*	10.09	2.71	6.25	2.00
4	3.91	3.37	2.03	2.01
5*	15.91	7.69	11.25	4.44
6	0.56	3.56	1.32	1.88
7	1.26	1.29	1.09	1.22
8*	11.12	3.46	7.19	2.26
9	3.29	2.83	1.70	1.76
Mean	7.52	3.30	4.96	2.15

These results confirm that integrating HSMM-based degradation modeling with particle filtering provides robust probabilistic RUL predictions. The flexibility of the PF enables the integration of prior knowledge about imperfect repair, which modifies the degradation trajectory at the time of repair. This adaptation is achieved by injecting the repair effectiveness prior into the particle ensemble. At the same time, uncertainty regarding the repair effect is explicitly propagated. As a result, repair uncertainty and prognostic uncertainty are treated in a unified Bayesian manner.

The advantage of the SBSM becomes more evident in specimens with thick repairs, which exhibit extended post-repair lifetimes beyond the range observed during training. While the CNN performs competitively for some thin repair cases (where post-repair lifetimes remain closer to the training distribution), it degrades in performance under larger distributional shifts. In contrast, the SBSM remains more consistent, reflecting its ability to modify degradation trajectories after repair without retraining.

The qualitative RUL trajectories further support these observations. Figure 5 presents the RUL prediction of specimen 7, which had a thin repair and the shortest post-repair life in the test set. After repair, the SBSM particles are shifted according to the prior $r \sim N(1.05, 0.02^2)$, and the prediction follows the true RUL closely. In this case, the CNN also

performs well, since the lifetime remains similar to those observed during training.




Figure 5. RUL prediction for specimen 7 (thin repair).

Figure 6 shows the RUL prediction of a specimen with a thick repair and the longest life in the test set. Due to the significant shift in the true RUL, the initial prior $r \sim N(1.3, 0.05^2)$ does not immediately match the new trajectory. However, as additional observations are incorporated, the SBSM progressively adjusts and moves closer to the true RUL. In contrast, the CNN cannot adequately adapt to this shift, resulting in poorer performance.

Overall, the results demonstrate that the proposed SBSM can provide probabilistic RUL predictions, incorporate imperfect repair in a principled Bayesian manner, and adapt to post-repair degradation dynamics without requiring repaired data during training. Compared to the CNN baseline, the SBSM shows improved robustness, better uncertainty calibration and sharpness, and more consistent performance across varying repair conditions.

5. CONCLUSIONS

This paper presents the Sequential Bayesian Semi Markov Framework (SBSM), which integrates HSMM as a degradation model with a PF to provide probabilistic RUL predictions




Figure 6. RUL prediction for specimen 5 (thick repair).

while explicitly accounting for imperfect repair. The degradation model was trained exclusively on non repaired specimens, and repair actions were incorporated during inference through prior distributions representing repair effectiveness. This formulation enables post-repair degradation dynamics to differ from pre-repair behavior without requiring repaired data during training.

Experimental validation on cold spray repaired aluminum fatigue specimens demonstrated that the proposed framework outperforms a CNN baseline in both point prediction accuracy (RMSE) and uncertainty calibration (CRPS). The results show that the SBSM provides more robust performance, particularly under significant distributional shifts induced by thick repairs. By injecting prior knowledge about repair effectiveness into the particle ensemble, the framework treats repair uncertainty and prognostic uncertainty in a unified Bayesian manner. This capability allows structured adaptation after repair and improves consistency across varying repair conditions.

Overall, the study confirms that combining stochastic degradation modeling with sequential Bayesian inference is effective for post-repair prognostics. The framework achieves probabilistic RUL prediction, principled integration of imperfect

repair, and adaptation to new degradation regimes without re-training.

Future work will focus on investigating the impact of different prior assumptions regarding repair effectiveness and analyzing how these priors influence RUL prediction accuracy and uncertainty. A more detailed uncertainty and sensitivity analysis will also be conducted to better understand the contribution of model parameters and prior distributions to predictive performance. Finally, the development of a real-time HI will enable the integration of the SBSM into a complete PHM pipeline suitable for fully real-time deployment.

ACKNOWLEDGMENTS

Part of this work was made possible by funding from the Dutch Research Council (NWO) Open Technology Programme, under grant number 20434 (CSAR).

REFERENCES

- ASTM International. (2021). *Standard practice for conducting force controlled constant amplitude axial fatigue tests of metallic materials* (No. E466-21). West Conshohocken, PA, USA: ASTM International. doi: 10.1520/E0466-21
- Carlo, F. D., & Arleo, M. A. (2017). Imperfect maintenance models, from theory to practice. In C. Volosencu (Ed.), *System reliability* (chap. 18). London: IntechOpen. doi: 10.5772/intechopen.69286
- Do, P., Voisin, A., Levrat, E., & Iung, B. (2015). A proactive condition-based maintenance strategy with both perfect and imperfect maintenance actions. *Reliability Engineering & System Safety*, 133, 22-32. doi: <https://doi.org/10.1016/j.ress.2014.08.011>
- Dong, M., & He, D. (2007). A segmental hidden semi-markov model (hsmm)-based diagnostics and prognostics framework and methodology. *Mechanical Systems and Signal Processing*, 21(5), 2248-2266. doi: <https://doi.org/10.1016/j.ymssp.2006.10.001>
- Gal, Y., & Ghahramani, Z. (2016). Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *international conference on machine learning* (pp. 1050–1059).
- Jouin, M., Gouriveau, R., Hissel, D., Péra, M.-C., & Zerhouni, N. (2016). Particle filter-based prognostics: Review, discussion and perspectives. *Mechanical Systems and Signal Processing*, 72-73, 2-31. doi: <https://doi.org/10.1016/j.ymssp.2015.11.008>
- Komninos, P., Galanopoulos, G., Kontogiannis, T., Eleftheroglou, N., & Zarouchas, D. (2025). A bayesian inference-based framework for modeling imperfect post-repair behavior of remaining useful life under uncertainty. *Expert Systems with Applications*, 288, 127723. doi: <https://doi.org/10.1016/j.eswa.2025.127723>
- Kontogiannis, T., Salinas-Camus, M., & Eleftheroglou, N. (2025). Chapter ten - hidden markov model applications: Aviation prognostics. In I. S. Triantafyllou, S. Malefaki, & A. Karagrigoriou (Eds.), *Stochastic modeling and statistical methods* (p. 191-213). Academic Press. doi: <https://doi.org/10.1016/B978-0-44-331694-4.00015-3>
- Kontogiannis, T., Salinas-Camus, M., & Eleftheroglou, N. (2026). Himap: Hidden markov for advanced prognostics. *Journal of Open Source Software*, 11(121), 9491.
- Ma, J., Cai, L., Liao, G., Yin, H., Si, X., & Zhang, P. (2023). A multi-phase wiener process-based degradation model with imperfect maintenance activities. *Reliability Engineering & System Safety*, 232, 109075. doi: <https://doi.org/10.1016/j.ress.2022.109075>
- Moleda, M., Malysiak-Mrozek, B., Ding, W., Sunderam, V., & Mrozek, D. (2023). From corrective to predictive maintenance—a review of maintenance approaches for the power industry. *Sensors*, 23(13). doi: 10.3390/s23135970
- Salinas-Camus, M., Goebel, K., & Eleftheroglou, N. (2025). A comprehensive review and evaluation framework for data-driven prognostics: Uncertainty, robustness, interpretability, and feasibility. *Mechanical Systems and Signal Processing*, 237, 113015. doi: <https://doi.org/10.1016/j.ymssp.2025.113015>
- Skordilis, E., & Moghaddass, R. (2020). A deep reinforcement learning approach for real-time sensor-driven decision making and predictive analytics. *Computers & Industrial Engineering*, 147, 106600. doi: <https://doi.org/10.1016/j.cie.2020.106600>
- Wang, N., Sun, S.-d., Cai, Z.-q., Zhang, S., & Saygin, C. (2014). A hidden semi-markov model with duration-dependent state transition probabilities for prognostics. *Mathematical Problems in Engineering*, 2014(1), 632702.
- Widener, C., Carter, M., Ozdemir, O., Hrabe, R., Hoiland, B., Stamey, T., ... Eden, T. J. (2016). Application of high-pressure cold spray for an internal bore repair of a navy valve actuator. *Journal of Thermal Spray Technology*, 25(1), 193–201.
- Yandouzi, M., Gaydos, S., Guo, D., Ghelichi, R., & Jodoin, B. (2014). Aircraft skin restoration and evaluation. *Journal of thermal spray technology*, 23(8), 1281–1290.
- Yu, S.-Z. (2010). Hidden semi-markov models. *Artificial Intelligence*, 174(2), 215-243. (Special Review Issue) doi: <https://doi.org/10.1016/j.artint.2009.11.011>