A Sequential Bayesian Semi-Markov Framework for Mission-Based Prognostics

Mariana Salinas-Camus[1] and Nick Eleftheroglou[2]

[1,2] *Intelligent System Prognostics for Operations and Maintenance Group, Department of Aerospace Structures and Materials, Faculty of Aerospace Engineering, Delft University of Technology, Kluyverweg 1, Delft, 2629HS, the Netherlands*
m.salinascamus@tudelft.nl
n.eleftheroglou@tudelft.nl

Abstract

Prognostics aims to predict the Remaining Useful Life (RUL) of engineering systems to support decision-making in Predictive Maintenance (PdM) and operational planning. While significant efforts have focused on maintenance optimization, the impact of operational decisions on system degradation remains relatively underexplored. In mission-based contexts, systems operate under varying operational profiles, which directly influence degradation dynamics and, consequently, RUL.

This paper proposes a Sequential Bayesian Semi-Markov (SBSM) framework for RUL prediction under changing operational profiles. System degradation is modeled using Hidden Semi-Markov Models (HSMMs) with profile-dependent sojourn time distributions and shared emission models, enabling consistent degradation representation across operatinal profiles. Online state estimation and uncertainty propagation are performed using a particle filter, allowing adaptive tracking of degradation under operational profile switching. A quantile-preserving transformation is introduced to ensure coherent state evolution during profile switching.

The framework supports two prediction modes: (i) a non-informed mode, assuming unknown future operation, and (ii) an informed mode, leveraging known future operational schedules for decision support. The approach is validated on a simulated battery degradation dataset with time-varying operational profiles. Results demonstrate that incorporating future operational information significantly improves both prediction accuracy and uncertainty calibration, while the proposed non-informed approach achieves competitive performance with improved interpretability compared to a CNN baseline.

These results highlight the potential of the SBSM framework for mission-based prognostics and its applicability to operational decision-making under uncertainty.

1. Introduction

Modern engineering systems are increasingly operated under mission-based paradigms, in which a system executes a sequence of missions characterized by changing operational profiles, objectives, and constraints (Wesendrup & Hellingrath, 2020). In this context, mission-based prognostics refers to the prediction of the Remaining Useful Life (RUL) of a system while explicitly accounting for the fact that system usage is driven by mission execution and may involve deliberate changes in operational profiles.

Prognostics is a key enabler of decision-making within Prognostics and Health Management (PHM) frameworks (Goebel et al., 2017). In mission-based settings, prognostic outputs are not only used to anticipate failures but also to inform operational post-prognostic decision-making, where operational profiles are adapted to balance performance, safety, and system longevity during mission execution (Wesendrup & Hellingrath, 2020). Mission-based operational decision-making—such as mission replanning or task allocation using RUL information—can be seen as a specific subset of operational post-prognostic decision-making.

Despite this potential, post-prognostic decision-making research has predominantly focused on maintenance actions. The review by (Bougacha et al., 2020) shows that nearly half of the existing studies address maintenance decisions, whereas only about a quarter consider operational decision-making. Consequently, the role of prognostic models specifically designed for mission-driven operational profile changes remains underexplored.

Several works have investigated operational decision-making using RUL information. Mission replanning based on prognostics was proposed for autonomous vehicles in (Tang et al., 2011), while (Zhang et al., 2014) developed an

RUL-enhanced mission planner for planetary rovers. Similar objectives were pursued in (Salinas-Camus et al., 2023) and (Tamssaouet et al., 2019). However, these approaches typically assume that an accurate prognostic model is already available and do not address how such models should be constructed to account for mission-induced changes in operational profiles.

Indeed, nearly half of post-prognostic decision-making studies do not specify how RUL is predicted (Bougacha et al., 2020), implicitly separating prognostics from decision-making. This separation overlooks their closed-loop interaction: operational decisions influence degradation, while prognostic predictions inform future decisions. From a prognostic perspective, this implies that models intended for mission-based applications must explicitly account for changing operational profiles and support sequential interaction with operational decision-making algorithms.

Another limitation of existing prognostic frameworks is the widespread use of deterministic RUL predictions (Salinas-Camus et al., 2025). Since RUL represents an uncertain prediction of future system behavior, uncertainty must be explicitly quantified. Moreover, prognostic models for mission-based applications must be interpretable, adaptive, and capable of incorporating known or unknown future operational schedules.

While most prognostic models assume fixed operating conditions, some works address time-varying operational profiles (Xue et al., 2025). Although deep learning approaches have been developed, they typically require large labeled datasets, offer limited interpretability, and produce deterministic predictions. For example, (Chang et al., 2022) proposed a multi-input Long Short-Term Memory (LSTM) model that incorporates future operational conditions. However, the resulting framework remains a black-box model and does not provide probabilistic RUL predictions. In contrast, stochastic and Bayesian approaches naturally support uncertainty propagation and can be trained in an unsupervised manner. Existing Bayesian models either neglect how operational profile changes affect the system lifetime distribution (Liu et al., 2024), rely on system-specific physics (Jain & Lad, 2020), or impose restrictive assumptions such as linearity or Gaussian noise (Li et al., 2019; Keizers et al., 2021). A notable exception is the work of (Flory et al., 2014), who proposed a switching diffusion model for systems operating in randomly varying environments. However, their approach focuses on lifetime estimation and does not explicitly quantify uncertainty in the resulting predictions. Consequently, there remains a need for a probabilistic prognostic framework capable of predicting RUL under changing operational profiles while rigorously propagating uncertainty.

The objective of this paper is to propose a prognostic model specifically tailored to mission-based operation, where operational profiles are deliberately changed to fulfill mission objectives. To this end, the Sequential Bayesian Semi-Markov (SBSM) framework is proposed for RUL prediction under changing operational profiles. Degradation is modeled using a Hidden Semi-Markov Model (HSMM) (Yu, 2010) with discrete degradation states and Weibull-distributed sojourn times. A finite set of operational profiles is assumed, with one HSMM trained per profile in an unsupervised manner. Sequential inference and RUL prediction are performed using a Particle Filter (PF), enabling non-linear and non-Gaussian uncertainty propagation (Jouin et al., 2016). The work most closely related to the proposed framework is that of (Xiao et al., 2018), who employed a duration-dependent HSMM combined with a high-order particle filter for online prognostics. In contrast, the proposed SBSM framework uses Weibull-distributed state durations and a standard PF. The Weibull distribution is widely adopted for modeling degradation processes in engineering systems (Pham, 2022), while the use of a standard PF reduces computational complexity and improves computational efficiency.

RUL prediction is carried out in two modes: (i) a non-informed mode, assuming no knowledge of future operational usage, and (ii) an informed mode, assuming the future mission schedule is known. By explicitly modeling mission-driven operational profile switching and uncertainty propagation, the proposed framework provides prognostic outputs suitable for use within mission-based operational post-prognostic decision-making.

The main contributions of this work are:

1. A prognostic framework tailored to mission-based operation with changing operational profiles.

2. An uncertainty-aware RUL prediction method based on HSMMs and particle filtering.

3. The formulation of informed and non-informed RUL prediction modes for mission-driven operational contexts.

The remainder of this paper is organized as follows. Section 2 presents the proposed methodology, including the SBSM degradation modeling, the online inference procedure, and the RUL prediction under both unknown and known future operational profiles. It also describes the benchmark CNN architecture used for comparison. Section 3 introduces the case study based on simulated batteries and details the parameters employed to generate the dataset. Section 4 reports the results of the SBSM, both informed and non-informed, alongside those of the CNN, evaluating prediction accuracy and uncertainty calibration. Finally, Section 5 summarizes the conclusions.

2. METHODOLOGY

The proposed Sequential Bayesian Semi-Markov (SBSM) framework combines offline degradation modeling with online Bayesian inference and RUL prediction. The framework consists of three main components: HSMM-based degradation modeling, particle-filter-based online state inference with operational profile adaptation, and RUL prediction under unknown or known future operational profiles.

2.1. HSMM-based degradation modeling

In the offline phase, historical condition monitoring (CM) data, or health indicators, are used to learn degradation models based on Hidden Semi-Markov Models (HSMMs) using the HiMAP package (Kontogiannis et al., 2026). An HSMM represents the degradation process as a sequence of discrete degradation states, where each degradation state corresponds to a distinct degradation level of the system. Each degradation state is associated with a state-dependent observation model and an explicit sojourn time distribution that governs the duration spent in that state.

Formally, an HSMM is defined by the number of degradation states N, the initial state distribution π, the emission distributions B, and the state sojourn time distributions λ. In this work, Gaussian emission distributions are considered, parameterized by a mean μ and standard deviation σ, while Weibull distributions are used to model state sojourn times, parameterized by a shape parameter α and a scale parameter β. The number of degradation states N is the only hyperparameter of the model. All remaining parameters are learned using the Expectation–Maximization algorithm and are denoted by $\theta = \{\pi, B, \lambda\}$.

For mission-based prognostics, a separate HSMM is trained for each operational profile using run-to-failure trajectories collected under that profile. It is assumed that all operational profiles share identical emission distributions B across states. This implies that the observed health indicator provides a consistent representation of the underlying degradation state (see, for example, the health indicator construction methodology for time-varying operational conditions proposed by (Yang et al., 2025)), independently of the operating conditions. Such an assumption is reasonable when a robust health indicator has been developed that accurately reflects the degradation process. Under this assumption, differences in degradation dynamics across operational profiles are captured exclusively through the state sojourn time distributions λ, which represent profile-dependent degradation rates.

Training provides a set of parameters θ for each operational profile, characterizing the corresponding degradation dynamics and serving as input to the online state inference and RUL prediction framework.

2.2. Online state inference via particle filtering

During online operation, the hidden degradation state is inferred sequentially using a PF (Jouin et al., 2016). Given observations $y_{1:k}$, the posterior distribution of the degradation state is approximated as

$$p(x_k \mid y_{1:k}) \approx \sum_{i=1}^{M} \omega_k^{(i)} \delta(x_k - x_k^{(i)}), \tag{1}$$

where each particle $x_k^{(i)} = (s_k^{(i)}, e_k^{(i)})$ represents the current degradation state and the elapsed time spent in that state.

Particles are propagated according to the HSMM transition model. State transitions are governed by the discrete-time Weibull hazard function

$$h_s(e) = 1 - \frac{S_s(e+1)}{S_s(e)}, \tag{2}$$

while particle weights are updated using the state-dependent observation likelihood

$$\omega_k^{(i)} \propto \omega_{k-1}^{(i)}\, p(y_k \mid s_k^{(i)}). \tag{3}$$

Systematic resampling is performed when particle degeneracy is detected using a threshold on the effective sample size.

2.3. Operational profile transition

Operational profile transitions are handled independently of the RUL prediction strategy and are applied whenever a change in the operational profile is detected.

Since a dedicated HSMM is trained for each operational profile, a profile change requires adapting the state duration model. Consider a particle in degradation state s with elapsed time e. Under the previous profile, the state duration follows a Weibull distribution with parameters $(\alpha_s^{old}, \beta_s^{old})$, while under the new profile it follows $(\alpha_s^{new}, \beta_s^{new})$.

To preserve the particle's relative progress within the degradation state, a quantile-preserving transformation is applied:

$$q = F_{old}(e), \qquad e' = F_{new}^{-1}(q), \tag{4}$$

where $F(\cdot)$ denotes the Weibull cumulative distribution function. This transformation ensures that the particle maintains the same percentile of the state duration distribution after the profile transition. The degradation state remains unchanged, and only the elapsed time is updated. Particles already in the absorbing failure state are not modified.

This quantile-preserving transformation is central to the proposed framework, as it enables consistent tracking of degradation progress across operational profile changes without resetting the state or losing information about the system's degradation history.

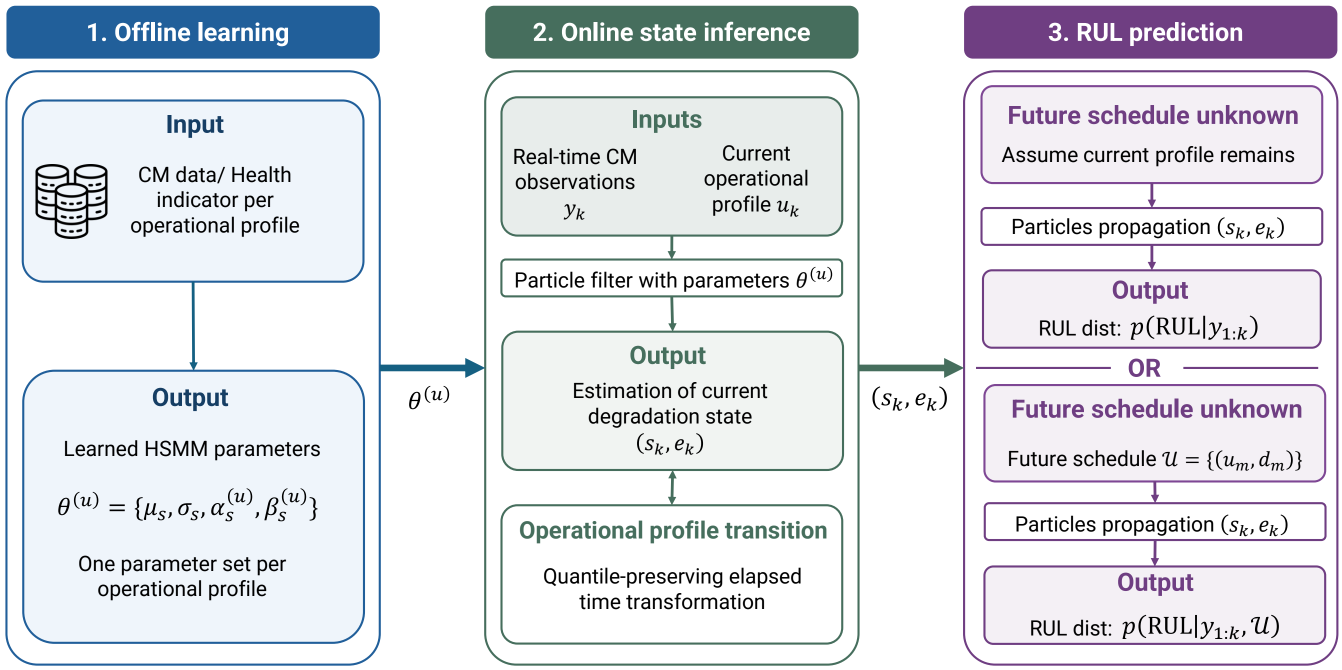

Figure 1. Overview of the SBSM framework for mission-based prognostics.

2.4. RUL prediction with unknown future operation

When future operational profiles are unknown, RUL prediction is performed by propagating the posterior particle ensemble forward in time while assuming that the current operational profile remains active.

For each particle, multiple Monte Carlo realizations of future degradation trajectories are generated. If a particle is already in the failure state, its RUL is set to zero. Otherwise, the remaining time in the current degradation state is sampled from a Weibull distribution conditioned on survival beyond the elapsed time:

$$T^{(i,j)} \sim \mathrm{Weibull}(\alpha_{s_k^{(i)}}, \beta_{s_k^{(i)}}), \quad T^{(i,j)} > e_k^{(i)}. \tag{5}$$

The particle is then propagated through successive degradation states until the absorbing failure state is reached.

The resulting RUL samples capture aleatoric uncertainty due to the stochastic nature of degradation, and epistemic uncertainty arising from uncertainty in the inferred degradation state. The predictive RUL distribution is obtained by aggregating all samples and weighting them according to the particle weights.

2.5. RUL prediction with known future operational schedule

When the future operational schedule is known, the degradation process is propagated through the prescribed sequence of operational profiles. Let

$$\mathcal{U} = \{(u_m, d_m)\}_{m=1}^{M} \tag{6}$$

denote the future operational profiles and their durations.

Starting from the posterior particle set, future degradation trajectories are simulated sequentially across the scheduled profiles. Within each profile, state transitions are governed by the corresponding Weibull hazard function

$$h_s(e, u_m) = 1 - \frac{S_s^{(m)}(e+1)}{S_s^{(m)}(e)}, \tag{7}$$

where the superscript (m) denotes the operational profile.

When the duration of an operational profile is exhausted, the current degradation state and elapsed time are preserved, while the Weibull parameters are updated to those of the next profile. Monte Carlo realizations are generated until the absorbing failure state is reached, providing a weighted approximation of the posterior predictive distribution

$$p(\mathrm{RUL} \mid y_{1:k}, \mathcal{U}). \tag{8}$$

By explicitly incorporating future operational scheduling, this formulation reduces uncertainty associated with unknown future usage and enables operationally informed prognostic predictions.

2.6. Benchmark model: Convolutional Neural Network

To evaluate the performance of the SBSM framework, a Convolutional Neural Network (CNN) is employed as the baseline model, given the widespread adoption of CNNs in prognostics following the emergence of deep learning techniques. The CNN is trained to map fixed-length temporal windows of CM data, together with a one-hot encoded representation of the corresponding operational profile, to the RUL. The network architecture is illustrated in Figure 2. The extracted feature representations are subsequently mapped to a scalar RUL prediction through a fully connected layer with a linear activation function. The network parameters θ are optimized by minimizing the mean squared error (MSE) between the predicted and true RUL values:

$$\mathcal{L}(\theta) = \frac{1}{N}\sum_{i=1}^{N}\left(y_i - f_\theta(\mathbf{X}_i)\right)^2, \tag{9}$$

where N denotes the number of training windows, y_i is the ground-truth RUL associated with the i-th input window, and $f_\theta(\mathbf{X}_i)$ represents the CNN prediction for input $\mathbf{X}_i$.

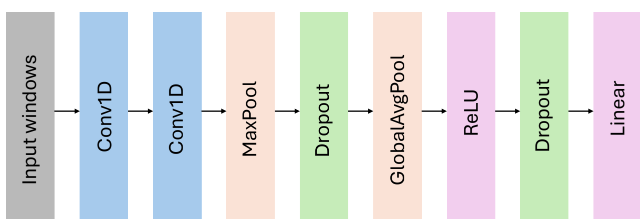

Figure 2. CNN architecture.

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 SBSM, the CNN does not explicitly model the system'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

The case study is based on a simulated battery degradation dataset generated using the ProgPy library developed by NASA (Teubert et al., 2023). Specifically, the dataset was generated using the `BatteryCircuit` model, with a process noise level of 0.01 and a measurement noise level of 0.03. The End of Life (EOL) is defined as the time at which the battery voltage reaches 3.1 [V].

Three operational profiles were defined based on the discharge current: Operation A with a current draw of 3.2 [A], Operation B with 3.5 [A], and Operation C with 3.7 [A].

For each operational profile, 20 run-to-failure degradation trajectories were generated, where the battery was operated under the same discharge current throughout its entire lifetime. These trajectories constitute the training dataset.

For testing, degradation trajectories with time-varying operational profiles were generated. Each battery experiences all three operational profiles exactly once during its lifetime, without repetition. The order of the profiles and the switching times are randomly sampled for each trajectory. In total, 20 degradation histories were generated for testing.

Figure 3 shows the EoL distribution for each dataset. It can be observed that, for the individual operational profiles (datasets A, B, and C), the EoL distributions are relatively concentrated. In contrast, the testing dataset exhibits a more dispersed distribution, reflecting the variability introduced by switching between operational profiles.

This experimental setup allows the evaluation of the proposed framework under controlled operational profile switches to validate the proposed prognostics framework capabilities.

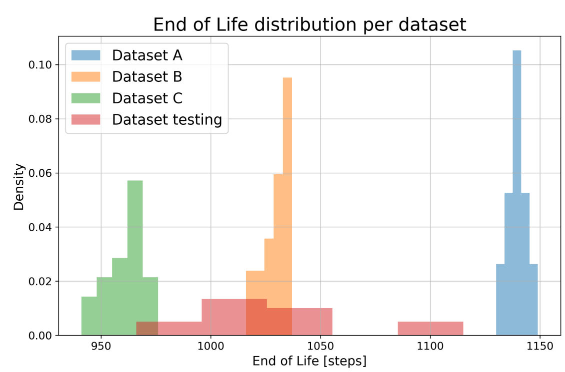

Figure 3. EOL distribution per dataset.

4. RESULTS

The SBSM framework requires training an HSMM. The Bayesian Information Criterion (BIC) was used to determine the optimal number of hidden states, resulting in a model with five states. A general HSMM was first trained using the entire pool of training data to estimate the shared emission parameters $\mathbf{B}$. These emission parameters were subsequently fixed and used to train three operation-profile-specific HSMMs. As all models share the same emission parameters, the key distinguishing factor between them lies in their sojourn time distributions, which capture the dynamics of each operational profile.

The learned parameters were then used within a particle filtering framework for online state inference, using 60 particles for both informed and non-informed SBSM variants.

For comparison, a CNN-based baseline model was implemented using a fixed input window length of 10. The model was trained with the Adam optimizer for up to 100 epochs, with early stopping applied to mitigate overfitting. Hyperparameters, including learning rate, batch size, and dropout rate, were optimized via random search, with the final model selected based on validation performance.

Model performance was evaluated using Root Mean Squared Error (RMSE) to assess point prediction accuracy, and Continuous Ranked Probability Score (CRPS) to jointly evaluate predictive accuracy, uncertainty calibration, and sharpness. The distributions of both metrics across trajectories are shown in Figure 4.

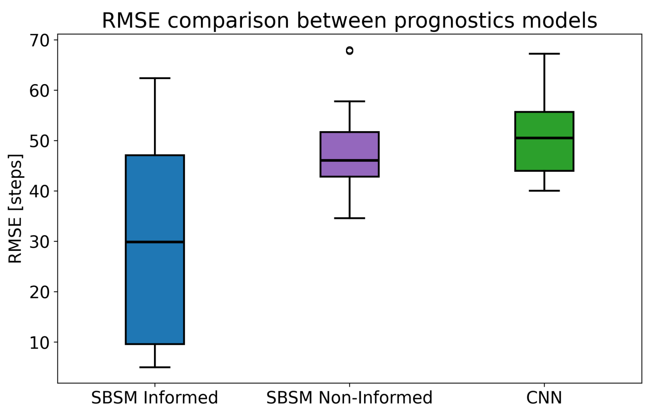

(a) RMSE results

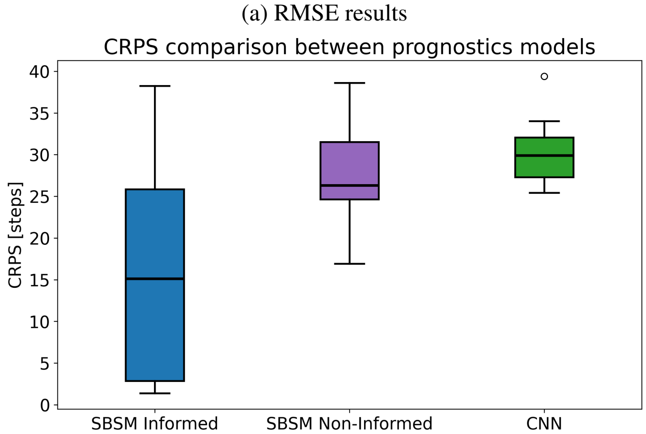

(b) CRPS results

Figure 4. Results in terms of prediction accuracy and uncertainty calibration for the three models.

The results indicate a clear performance advantage for the SBSM-informed approach. In terms of RMSE (Figure 4a), the SBSM-informed model achieves the lowest median error (29.88 steps), significantly outperforming both the SBSM non-informed model (median of 46.07 steps) and the CNN baseline (median of 50.51 steps). Additionally, the SBSM-informed model exhibits a wider spread with lower minimum values, suggesting that it can achieve highly accurate predictions under favorable conditions, although with some variability.

A similar trend is observed for CRPS (Figure 4b), where the SBSM-informed model again substantially outperforms the alternatives. The median CRPS is 15.12 steps, compared to around 26.30 steps for the SBSM non-informed model and 29.87 steps for the CNN. This improvement highlights the benefit of incorporating future operational profile information, enabling better-calibrated predictive distributions with improved sharpness.

Overall, these results demonstrate that incorporating knowledge of future operational profiles significantly enhances both prediction accuracy and uncertainty quantification. This is expected, as only the SBSM-informed model has access to future operational schedules. Therefore, the most meaningful comparison is between the SBSM non-informed model and the CNN baseline, which operate under the same information constraints. Under these conditions, the two models show comparable performance in both RMSE and CRPS, with the SBSM non-informed model achieving slightly better median values.

Beyond predictive performance, the SBSM framework offers substantial advantages in interpretability. By combining an HSMM-based degradation model with a particle filter for online inference, the evolution of the degradation state can be tracked explicitly. Each hidden state corresponds to a distinct degradation level, which can be visualized and analyzed over time. Furthermore, the predictive log-likelihood of observations can be monitored, providing a quantitative measure of how well the current observations align with the estimated degradation state. This enables users to diagnose model inconsistencies: for instance, if the estimated degradation state is incorrect, it is likely to propagate into inaccurate RUL predictions.

In contrast, the CNN model operates as a direct mapping from input data (condition monitoring signals and operational profiles) to RUL predictions. Due to its high-dimensional parameterization, the model behaves as a black box, offering limited insight into the underlying degradation process or the cause of prediction errors.

Additionally, the two approaches differ in their treatment of uncertainty. The CNN relies on MC dropout to approximate epistemic uncertainty, whereas the SBSM framework provides a more comprehensive uncertainty representation. Specifically, aleatoric uncertainty is inherently captured through the HSMM degradation model, while epistemic uncertainty is accounted for through the particle filtering process. This results in a richer and more interpretable characterization of predictive uncertainty.

As a qualitative example, Figure 5 presents the RUL prediction for trajectory 18. The predictions of the CNN and the SBSM non-informed model are broadly similar, although the CNN exhibits higher variability and less smooth predictions. In contrast, the SBSM-informed model produces a more stable prediction with reduced uncertainty bounds, reflecting the additional information about the future operational schedule. This demonstrates the practical advantage of leveraging operational profile schedule in prognostic modeling.

Overall, the SBSM framework can be leveraged for mission-based applications in two complementary ways. First, in scenarios where operational decisions dynamically alter the system usage, the framework provides reliable RUL predictions even when the future operational schedule is unknown. Second, when the decision-maker seeks to evaluate alternative operational strategies, the informed SBSM mode enables the comparison of different future operational schedules. This allows assessing how various mission plans would impact the system's degradation and RUL, thereby supporting decision-making processes aimed at optimizing objectives such as system longevity, safety, or performance.

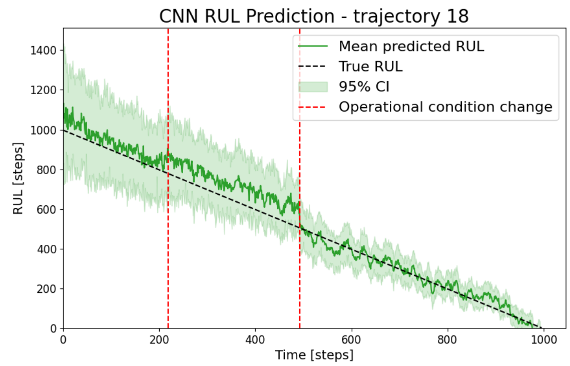

(a) CNN.

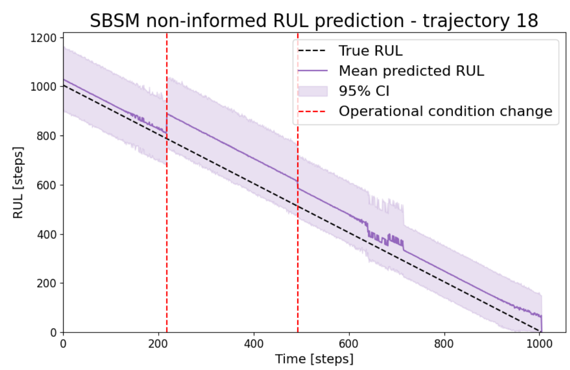

(b) SBSM non-informed.

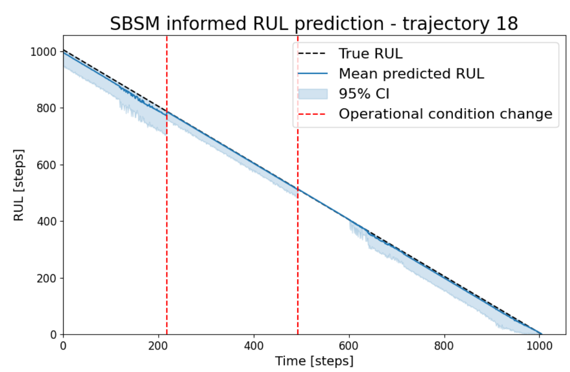

(c) SBSM informed.

Figure 5. RUL prediction for an example trajectory under different models.

5. CONCLUSIONS

This paper proposed a Sequential Bayesian Semi-Markov (SBSM) framework for mission-based prognostics under changing operational profiles. The approach combines HSMM-based degradation modeling with particle filtering for online inference, enabling both accurate RUL prediction and principled uncertainty quantification in dynamic operational environments.

The proposed framework explicitly accounts for changes in operational profiles by associating each profile with a specific HSMM, while sharing emission parameters across models to ensure consistency in degradation representation. Two RUL prediction modes were introduced: a non-informed mode, which assumes unknown future operation, and an informed mode, which leverages knowledge of future mission schedules. This distinction allows the framework to be applied in both real-time operational contexts and decision-support scenarios.

The results demonstrate that incorporating future operational information significantly improves predictive performance. The SBSM-informed approach achieved the lowest RMSE and CRPS values, indicating both high accuracy and well-calibrated probabilistic predictions. Furthermore, the proposed framework enables uncertainty management specifically for future uncertainties, defined as the reduction of predictive uncertainty by leveraging knowledge of future operational conditions. When the future operation is known, the model can condition its predictions on the corresponding operational profile, leading to narrower confidence intervals in RUL estimates. Under comparable information conditions, the SBSM model without future information and the CNN baseline exhibited similar performance, with the SBSM showing slightly better median results.

Beyond predictive accuracy, the SBSM framework provides important advantages in interpretability and uncertainty representation. Unlike the CNN baseline, which operates as a black-box mapping, the SBSM enables explicit tracking of degradation states and their evolution over time. Furthermore, the framework naturally captures both aleatoric uncertainty, through the stochastic degradation model, and epistemic uncertainty, through particle-based inference, resulting in richer and more interpretable predictive distributions.

The proposed framework also supports mission-based decision-making by enabling the evaluation of alternative operational strategies. By simulating different future operational schedules, decision-makers can assess their impact on system degradation and RUL, thereby supporting objectives such as extending system lifetime or ensuring mission completion.

While the SBSM framework has potential to be used in mission-based contexts and operational post-prognostics decision making, future work will focus on extending the framework to handle continuous and potentially unknown operational profiles. Further research will also explore integration with decision-making algorithms in closed-loop settings, enabling joint optimization of prognostics and operational planning. Finally, validation on real industrial datasets will be pursued to assess the practical applicability and robustness of the proposed approach.

REFERENCES

Bougacha, O., Varnier, C., & Zerhouni, N. (2020). A Review of Post-Prognostics Decision-Making in Prognostics and Health Management. *International Journal of Prognostics and Health Management*, *11*(15), 31. Retrieved from `https://hal.science/hal-03107810`

Chang, L., Lin, Y.-H., & Zio, E. (2022). Remaining useful life prediction for complex systems considering varying future operational conditions. *Quality and Reliability Engineering International*, *38*(1), 516-531. Retrieved from `https://onlinelibrary.wiley.com/doi/abs/10.1002/qre.2997` doi: https://doi.org/10.1002/qre.2997

Flory, J. A., Kharoufeh, J. P., & Gebraeel, N. Z. (2014). A switching diffusion model for lifetime estimation in randomly varying environments. *IIE Transactions*, *46*(11), 1227–1241. Retrieved from `https://doi.org/10.1080/0740817X.2014.893400` doi: 10.1080/0740817X.2014.893400

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).

Goebel, K., Daigle, M., Saxena, A., Sankararaman, S., Roychoudhury, I., & Celaya, J. (2017). *Prognostics: The science of making predictions*. CreateSpace Independent Publishing Platform. Retrieved from `https://books.google.nl/books?id=M0AztAEACAAJ`

Jain, A. K., & Lad, B. K. (2020). Prognosticating ruls while exploiting the future characteristics of operating profiles. *Reliability Engineering & System Safety*, *202*, 107031. doi: https://doi.org/10.1016/j.ress.2020.107031

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

Keizers, L., Loendersloot, R., & Tinga, T. (2021). Unscented kalman filtering for prognostics under varying operational and environmental conditions. *International Journal of Prognostics and Health Management*, *12*(2).

Kontogiannis, T., Salinas-Camus, M., & Eleftheroglou, N. (2026). Himap: Hidden markov for advanced prognostics. *Journal of Open Source Software*, *11*(121), 9491.

Li, N., Gebraeel, N., Lei, Y., Bian, L., & Si, X. (2019). Remaining useful life prediction of machinery under time-varying operating conditions based on a two-factor state-space model. *Reliability Engineering & System Safety*, *186*, 88-100. doi: https://doi.org/10.1016/j.ress.2019.02.017

Liu, X., Lunde, E., Diehl, F., Zhang, A., Vatn, J., & Yin, S. (2024). Degradation modeling under time-varying operating conditions: Inference and prognosis with particle filter. *Reliability Engineering & System Safety*, *245*, 109965. doi: https://doi.org/10.1016/j.ress.2024.109965

Pham, H. (2022). Distribution functions and its applications. In *Statistical reliability engineering: Methods, models and applications* (pp. 67–134). Cham: Springer International Publishing. Retrieved from `https://doi.org/10.1007/978-3-030-76904-8_2` doi: 10.1007/978-3-030-76904-8_2

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: 10.1016/j.ymssp.2025.113015

Salinas-Camus, M., Kulkarni, C. S., & Orchard, M. E. (2023). Battery state-of-health aware path planning for a mars rover. In *15th annual conference of prognostics and health management.*

Tamssaouet, F., Nguyen, K., & Medjaher, K. (2019). System remaining useful imaximization through mission profile optimization. In *Asia-pacific conference of the prognostics and health management society.*

Tang, L., Hettler, E., Zhang, B., & DeCastro, J. (2011). A testbed for real-time autonomous vehicle phm and contingency management applications. In *Annual conference of the phm society* (Vol. 3).

Teubert, C., Jarvis, K., Corbetta, M., Kulkarni, C., & Daigle, M. (2023). Progpy: Python packages for prognostics and health management of engineering systems. *Journal of Open Source Software*, *8*(87), 5099. Retrieved from https://doi.org/10.21105/joss.05099 doi: 10.21105/joss.05099

Wesendrup, K., & Hellingrath, B. (2020). A process-based review of post-prognostics decision-making. In *Phm society european conference* (Vol. 5, pp. 12–12).

Xiao, Q., Fang, Y., Liu, Q., & Zhou, S. (2018). Online machine health prognostics based on modified duration-dependent hidden semi-markov model and high-order particle filtering. *The International Journal of Advanced Manufacturing Technology*, *94*(1), 1283–1297. Retrieved from https://doi.org/10.1007/s00170-017-0916-7 doi: 10.1007/s00170-017-0916-7

Xue, B., Yu, Y., Fu, H., Zhu, J., & Xu, Z. (2025). Advanced prognostics techniques for machinery under time-varying operating conditions: a review. *Applied Intelligence*, *55*(16), 1066.

Yang, J., Zhou, J., Chai, Y., Chen, D., & Qin, Y. (2025). Benchmark transformation neural network for health indicator construction under time-varying speed and its application in machinery prognostics. *Reliability Engineering & System Safety*, *257*, 110823. Retrieved from https://www.sciencedirect.com/science/article/pii/S0951832025000262 doi: https://doi.org/10.1016/j.ress.2025.110823

Yu, S.-Z. (2010). Hidden semi-markov models. *Artificial Intelligence*, *174*(2), 215-243. (Special Review Issue) doi: 10.1016/j.artint.2009.11.011

Zhang, B., Tang, L., DeCastro, J., Roemer, M., & Goebel, K. (2014). Autonomous vehicle battery state-of-charge prognostics enhanced mission planning. *International Journal of Prognostics and Health Management*, *5*(2).