Manufacturing Quality–Informed Prognostics: A Novel Approach to Past Uncertainty Management

Benjamin Brito Schiele[1], Julie Teuwen[2], and Nick Eleftheroglou[1]

[1] *Intelligent System Prognostics Group, Aerospace Structures and Materials Department, Faculty of Aerospace Engineering, Delft University of Technology*
b.a.britoschiele@tudelft.nl
n.eleftheroglou@tudelft.nl

[2] *Aerospace Structures and Materials Department, Faculty of Aerospace Engineering, Delft University of Technology*
j.j.e.teuwen@tudelft.nl

ABSTRACT

The future behavior of a system is largely determined by its manufacturing process, as variations in production quality can lead to different performance outcomes over time. In the context of prognostics, all sources of uncertainty that exist prior to the system's commencement of operation are collectively referred to as past uncertainty, yet its role is rarely recognised in the prognostics and health management (PHM) community. Most existing approaches to uncertainty and its management focus only on model parameters, leaving the management of past uncertainty largely unexplored. This work introduces a framework to explicitly manage this source by incorporating manufacturing quality control (MQC) data into hidden semi-Markov model (HSMM) prognostics. The method creates quality-specific HSMMs, each tailored to a particular manufacturing quality (MQ) type, and combines them during inference using MQC-informed Bayesian model averaging. The framework is validated on composite specimens with pristine, oil-induced, and Teflon-induced defects. Ultrasonic scans provide MQC inputs, while strain data describes degradation. Results show that accounting for MQ reduces uncertainty in RUL predictions and highlights the importance of correctly identifying the MQ type for effective uncertainty management in prognostics.

1. INTRODUCTION

Prognostics is a key component of prognostics and health management (PHM) solutions. In this context, the goal of prognostics is to predict the remaining useful life (RUL) of the system under study, which will later be used for decision-making. An appropriate implementation of prognostics enables cost savings and increased service availability (Goli, Ghodrati, & Eleftheroglou, 2025).

Among the many factors that concern prognostics, the inclusion of uncertainty is essential, since they aim to predict future events (Sankararaman & Goebel, 2013). To understand the origin of prognostic uncertainty, Sankararaman (Sankararaman, 2015) proposes four sources of uncertainty:

- Present uncertainty reflects the lack of knowledge about the current damage condition.
- Future uncertainty group uncertainties related to the future operational conditions that the system will face.
- Modelling uncertainty corresponds to uncertainty in the prognostic model parameters.
- Prediction uncertainty accounts for uncertainty generated due to approximations in the computation of the prognostic.

Later, (Eleftheroglou, 2020) introduced a fifth source of uncertainty: past uncertainty. This additional source refers to all uncertainties present before a system begins operation, such as those associated with its materials, assembly processes, and manufacturing quality.

Prognostic models must account for all five sources of uncertainty to provide a reliable estimate of the RUL uncertainty. Neglecting any of these sources can result in incomplete RUL uncertainty estimates, which may in turn mislead subsequent decision-making. Nevertheless, even when all sources are considered, the RUL uncertainty may remain so large that the prognostics output provides limited or no actionable information for decision-making. In these cases, it is necessary to obtain additional information to reduce uncertainty at spe-

cific sources, thereby reducing prognostic uncertainty before making any decisions. This process is referred to as uncertainty management (UM) (Sankararaman, 2015).

Some UM frameworks are available in PHM literature. For instance, (Saha & Goebel, 2008) explore the use of Bayesian techniques to manage present and modelling sources of uncertainty. In a similar fashion, (Orchard, Kacprzynski, Goebel, Saha, & Vachtsevanos, 2008) employ outer feedback correction loops to manage present, modelling, and prediction uncertainties. Furthermore, (Edwards, Orchard, Tang, Goebel, & Vachtsevanos, 2010) implements this technique for future uncertainty management with favourable results. (Tang, Kacprzynski, Goebel, & Vachtsevanos, 2009) review several techniques for imaging present and modelling sources of uncertainty in particle filter-based prognostics. Recently, (Kim, Choi, & Kim, 2022) proposed a framework for modelling uncertainty management with variable inspection schedules. Also, (Eleftheroglou, Zarouchas, & Benedictus, 2020; Eleftheroglou, Galanopoulos, & Loutas, 2024) have proposed frameworks for present uncertainty management using adaptive hidden semi-Markov models (HSMMs).

From the reviewed literature, it is evident that most UM frameworks focus primarily on managing present and modelling sources of uncertainty.

UM for past state uncertainty has, so far, received little attention. Initial results were presented by Leinarts (Leinarts, 2025), which proposes a similarity-based approach to adapt HSMM parameters using manufacturing quality control (MQC) data from composite specimens. However, it does not formalize a framework for managing past uncertainty in general PHM solutions.

Therefore, this work constitutes a first attempt to formalize a framework for managing past uncertainty. The proposed framework can use MQC data to better characterize the system's quality before its operation begins and use this information to update the prognostic model, resulting in more reliable RUL prognostics.

The paper is structured as follows. First, Section 2 presents the proposed framework. Then, Section 3 introduces the case study used for validating the framework. Later, the obtained results are presented and discussed in section 4. Finally, Section 5 concludes this work and provides future research directions.

2. PROPOSED FRAMEWORK

2.1. Assumptions

Past uncertainty management involves gathering additional information to assess the system's quality prior to operation and incorporating it into the chosen prognostic model before starting operation. The proposed framework achieves this goal by using MQC data to adjust the prognostic model's parameters related to the system's degradation dynamics.

As a starting point, this study assumes that the prognostic model corresponds to an HSMM, since this model offers several advantages for modelling prognostic uncertainty compared to other common approaches, such as neural network-based methods (Salinas-Camus & Eleftheroglou, 2024). These advantages lead to enhanced robustness, interpretability, and feasibility (Salinas-Camus, Goebel, & Eleftheroglou, 2025).

More specifically, HSMMs infer a hidden stochastic process that models the evolution of system degradation. To accomplish this, these models employ an observation process, a monotonic degradation process, and a distribution for the initial condition. The observation process, parameterized by B, infers the hidden damage state from condition monitoring (CM) data. Then, the degradation process, represented by Γ, models the left-to-right transition between hidden states. Finally, the initial state probabilities, parameterized by π, indicate the system's initial damage condition. The values for these parameters are estimated by maximizing the observation likelihood, as presented in Eq. 1:

$$\left\{\pi^*, B^*, \Gamma^*\right\} = \arg\max_{\pi,B,\Gamma} \sum_{i=1}^{K} log\big(Pr(y^{(k)}|\pi, B, \Gamma)\big)\ , \quad (1)$$

where y^k corresponds to the CM data of the k-th degradation history in a training dataset, $\left\{\pi^*, B^*, \Gamma^*\right\}$ denote the estimated parameters of the HSMM. In addition, $Pr(y^{(k)}|\pi, B, \Gamma)$ the likelihood function to be maximized. For further details on the mathematical formulation behind HSMMs, please refer to (Kontogiannis, Salinas-Camus, & Eleftheroglou, 2025).

The proposed framework assumes that sources of past uncertainty influence the dynamics of system degradation and its initial state. This implies that the observed CM data may remain the same, while the system's degradation rate changes with past uncertainty. In the HSMM, this assumption implies that the parameter B is independent of the sources of past uncertainty, while the π and Γ are conditional on them. Thus, managing past uncertainty entails using MQC data to update these parameters and reduce uncertainty in the output prognostics.

Another core assumption of this framework is that past uncertainty can be described by a discrete label representing the overall system manufacturing quality (MQ). This means that past uncertainty originates from the lack of knowledge about which MQ type the system belongs to. Note that MQ reflects the phenomena that induce past uncertainty without loss of generality.

Finally, three distinct data sources are necessary to implement the proposed framework. First, a dataset comprising example

degradation histories with CM and MQC data for all identified MQ types. The purpose of this dataset is to adjust the models that compose this framework. Second, before starting the operation of the system, a user should gather MQC data using a non-destructive testing technique. This data should reflect the system's MQ type and will later be used to perform past UM. Finally, during the operation, CM data is continuously collected and used as input to generate RUL prognostics.

2.2. Proposed framework

The proposed framework comprises two phases: training and inference. First, the training phase uses example degradation histories to train a set of HSMMs. These models reflect the system's degradation process, conditional on the detection of a particular MQ type, and are denoted as Quality-specific HSMMs (Q-HSMMs).

Then, the inference phase uses the Q-HSMMs to manage past uncertainty and predict the system's RUL. This is achieved by using MQC data obtained before the start of the operation to classify the MQ type. Then, the obtained classification is used to integrate the RUL prognostics of the Q-HSMMs using Bayesian model averaging. Note that each Q-HSMM uses the CM data gathered during operation as input for its RUL prognostics.

Figure 1 presents a flowchart of the proposed framework and the submodules it employs. Subsection 2.2.1 provides further details about the training phase of the framework, while Subsection 2.2.2 explains the inference phase.

2.2.1. Training phase

The goal of the training phase is to generate a collection of Q-HSMM using the available training degradation histories. The first assumption of this framework is that past sources of uncertainty affect only the degradation dynamics and the initial damage condition of the system. This means that all Q-HSMMs should share a common observation process. To achieve this, the models are trained in two steps.

In the first step, it is necessary to train a baseline HSMM that does not distinguish between MQ types. The purpose of this model is to infer the general observation process, parameterized by B, that relates the CM data to the hidden states. This is achieved by maximizing the observation likelihood as stated in Eq. 1.

Then, in the second step, the π and Γ parameters of the Q-HSMMs are estimated using subsets of the training dataset composed of only one MQ type. Note that for this step, the B parameter is not further updated. By doing this, the Q-HSMM is adjusted to the degradation dynamics induced by each MQ type without altering the observation process. Formally, this estimation procedure can be expressed in terms of the maximum likelihood estimation problem stated in Eq. 2:

$$\left\{\pi_i^*, \Gamma_i^*\right\} = \arg\max_{\pi,\Gamma} \sum_{k=1}^{K_i} log\big(Pr(x^{(k)}|\pi, \Gamma, B)\big)\,, \qquad (2)$$

where $\left\{\pi_i^*, \Gamma_i^*\right\}$ correspond to the estimated parameters of the i-th model, $x^{(k)}$ corresponds to the CM data of the k-th degradation history of the training dataset with MQ type i and size K_i. Additionally, B represents the observation process of the baseline HSMM. This optimization problem can be solved iteratively using the expectation-maximization algorithm (Eleftheroglou et al., 2020).

The outcome of repeating this procedure for all MQ type results in the desired collection of Q-HSMMs. Each of these models reflects the degradation dynamics caused by one specific MQ type, while sharing one common observation process. Because of this property, the Q-HSMMs isolate the effects of past uncertainties, resulting in more reliable prognostic models when used during inference.

2.2.2. Inference phase

At the inference phase, the Q-HSMMs generated during the training phase are used to predict the RUL of a new system. To do so, before the operation begins, it is necessary to leverage MQC data from the new system to classify its MQ type. This task involves computing the probability that the system exhibits a specific MQ type. This classification is later used to determine which Q-HSMM best represents the system under study.

Once the operation starts and CM data is collected, each Q-HSMM can generate its own RUL prognostics. Then, the effect of past uncertainties is marginalized by performing Bayesian model averaging, using the classification outcome as weights. Formally, this corresponds to computing the expression presented in Eq. 3:

$$Pr(RUL|x, q) = \sum_{i=1}^{N} Pr(RUL|x, \pi_i^*, \Gamma_i^*, B) \cdot Pr(\pi_i^*, \Gamma_i^*|q), \qquad (3)$$

where $(RUL|x, q)$ corresponds to the RUL prognostic probability density function (PDF) given the input CM data x and the MQC data q, $Pr(RUL|x, \pi_i^*, \Gamma_i^*, B)$ the RUL prognostic PDF for the i-th Q-HSMM and CM data x, $Pr(\pi_i^*, \Gamma_i^*|q)$ the output of the MQ type classifier given the MQC data q, and N the number of different MQ types available. Note that for the case of continuous MQ type, the summation in Eq. 3 can be generalized to a Lebesgue integral.

The RUL prognostic obtained from Eq. 3 factors in information from both CM and MQC data. However, including the classifier outputs in the RUL prognostic computation indi-

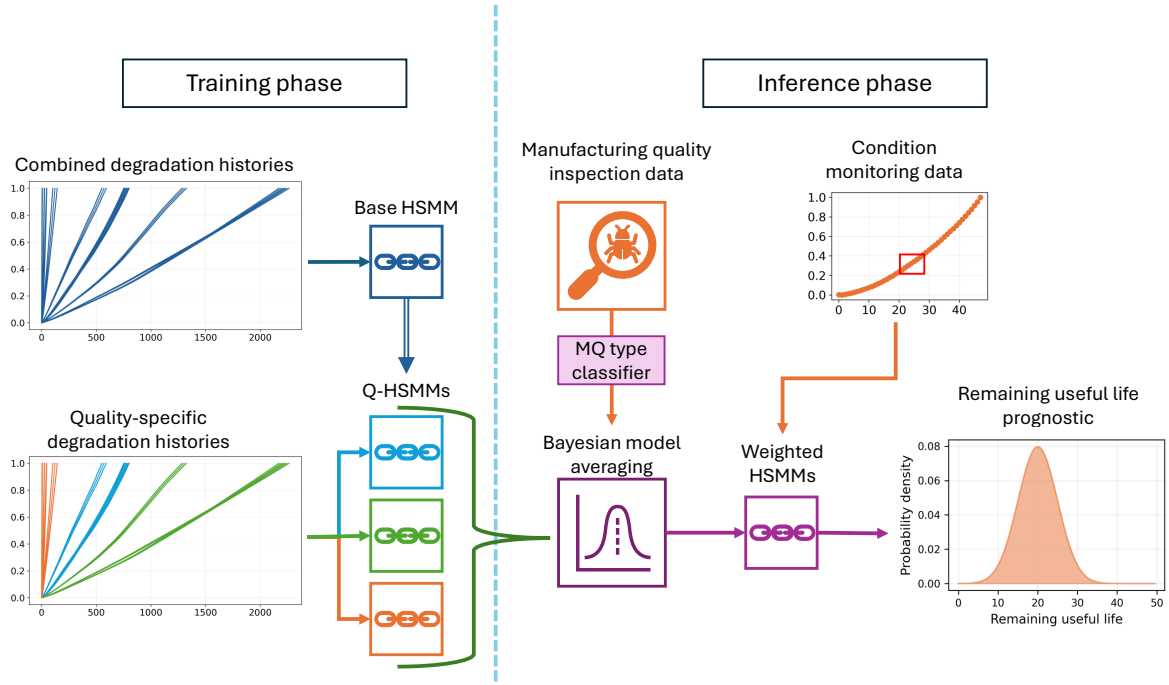

Figure 1. Flowchart of the proposed framework. Dark blue icons represent models trained with the combined degradation histories. Blue, green, and orange icons correspond to modules and data for a particular MQ type. Violet icons represent Q-HSMMs after combining them using Bayesian model averaging.

cates that the performance of the combined prognostic model depends not only on the individual Q-HSMMs but also on the classifier. If the classifier assigns a high probability to the correct MQ type, UM is effectively applied, since the effect of past uncertainty would be drastically reduced. Yet if the classifier is biased towards one particular MQ type or assigns uniform probabilities to all MQ types, it will systematically assign low probabilities to the correct Q-HSMM, leading to biased and uncertain prognostic outputs.

To isolate this effect of the classifier and showcase the potential of correctly applying the proposed uncertainty management framework, this work implements an ideal classifier to select the prognostic model given some MQC data. In the ideal case, the classifier assigns a probability of one to the correct MQ type and zero to any other. So, using the ideal classifier to generate prognostics yields the best possible performance, given a set of Q-HSMMs. This variant of the framework is referred to as ideal Q-HSMM.

However, achieving an ideal classifier is not always possible, since MQC data is also prone to observation noise and modelling uncertainty. To represent this realistic scenario, this work also implements a k-nearest neighbour (KNN) classification approach. This approach was chosen for its simplicity and its widespread use in classification tasks (Syriopoulos, Kalampalikis, Kotsiantis, & Vrahatis, 2025).

The implemented KNN classifier assigns MQ type probabilities based on the labels of the k nearest training examples. Distances are computed by projecting the MQC data into an embedding space and measuring the 2-norm between the projection and the training samples.This variant of the framework is referred to as KNN Q-HSMM.

Note that this implementation reflects only a simple, realistic classification approach. This work will not focus on achieving the best possible performance in this aspect.

To validate the proposed uncertainty management framework, the following sections implement and evaluate it using data from an experimental campaign.

3. CASE STUDY

The experimental campaign reported in (Leinarts, 2025) was selected as a case study to validate the proposed framework, as it contains all the necessary information sources to implement it. This dataset consists of 12 constant-amplitude fatigue tests on open-hole carbon fiber composite specimens. The fatigue test stop criterion is met when the specimen ruptures. During specimen manufacturing, two types of defects were deliberately introduced, resulting in three MQ types: pristine specimens, oil-induced defects, and Teflon-induced defects. Hence, past uncertainties arise from the quality of the manufacturing process and may include defects such as oil contamination or Teflon residues.

Before performing the fatigue test, all specimens were scanned

using ultrasonic imaging (Endrerud, 2014). These images reveal the quality of the laminate by showing damping information, highlighting air/dirt inclusions, delaminations, or voids. Thus, they can partially indicate which defects were introduced into the tested specimens. Thus, they serve as the MQC data in the proposed framework.

Figure 2 illustrates the MQC data for the available MQ types. In pristine specimens, scans still present noise. However, they do not reveal major defects. In contrast, specimens with Teflon-induced defects yield scans with square structures that correspond to the presence of Teflon patches. Finally, specimens with oil-induced defects tend to reveal elevated noise levels.

During the fatigue test, strain data were collected using digital image correlation (DIC). This data reflects the degradation of each specimen during the experiments. Figure 3 illustrates the raw time series data. The dataset was augmented by randomly compressing and stretching the DIC trajectories and adding white noise. Then, to isolate the effect of measurement noise and data quality, the augmented trajectories were further preprocessed to generate a cumulative normalized health indicator (HI) that serves as CM data in the proposed framework.

Figure 4a illustrates the degradation histories obtained after augmentation and preprocessing, with colors indicating the respective MQ type. Figure 4b reveals the impact of past uncertainties on the lifetime distributions of the systems. Indeed, separating the distributions by color, i.e., knowing each specimen's MQ type, leads to less spread in the lifetime distributions.

To quantify the performance of the proposed framework, two evaluation metrics are employed: the mean absolute error (MAE) and the continuous ranked probability score (CRPS). The MAE quantifies the deviation of the expected value of the prognostic PDF from the actual RUL curve, i.e., its bias. Then, the CRPS simultaneously evaluates the calibration of the uncertainty in the prognostic model and its bias. Hence, both the MAE and CRPS will quantify the performance of the model in terms of accuracy and uncertainty.

The HI degradation histories presented in Figure 4a. were used to train three Q-HSMMs, one for each MQ type, i.e., while leaving three degradation histories as a holdout test set. In addition, the scans were used to train a KNN MQ type classifier. With these elements, it is possible to implement the proposed framework and evaluate its performance.

4. RESULTS

The results obtained compare cases in which the KNN and ideal classifiers are used. The KNN classification will represent a realistic case in which uncertainty is managed using MQC data. In contrast, using the ideal classifier showcases the best possible performance that the framework can achieve in this specific case study. Additionally, results are contrasted with a baseline HSMM, representing the case in which past uncertainty is not managed, and only CM data is available.

Figure 5 shows the evolution of RUL prognostics across time for a hold-out pristine specimen for the baseline HSMM, the ideal, and the KNN Q-HSMMs. Among the three tested models, the ideal Q-HSMM has the lowest uncertainty, followed by the KNN Q-HSMM. Furthermore, when looking at the expected values, all of the obtained prognostics are close to the Actual RUL. This indicates that the Q-HSMMs can manage past uncertainty without introducing significant bias into the model.

Similar to the pristine specimens, Figure 6 presents the obtained prognostics for a hold-out specimen with a Teflon-induced defect. In this case, it can be observed that the ideal Q-HSMM manages past uncertainty. However, the obtained prognostics present a bias towards higher RUL values. This bias can be explained by the dispersion in lifetimes of specimens with Teflon-induced defect. Since these specimens are biased towards higher lifetimes, their associated Q-HSMM will inherit this bias.

The KNN Q-HSMM is, however, unable to manage past uncertainty in this particular specimen. This result can be attributed solely to the KNN MQ type classifier, since the ideal Q-HSMM does not exhibit this issue. In fact, for this specimen, the classifier assigns an equal probability to all Q-HSMMs, resulting in an equally weighted combination of all quality-specific prognostic models. Ultimately, this results in prognostics that still present excessive uncertainty.

The previous results highlight the importance of the MQ type classifier for managing past uncertainty. Indeed, if the classifier is biased towards a specific MQ type or too uncertain in its outputs, using it for past uncertainty management can result in worse performance than the baseline model. Hence, understanding which requirements should be enforced on the classifier to guarantee effective past UM constitutes an unattended research gap.

The previous analysis is based on results from a single dataset partition. To make the analysis robust to dataset selection bias, cross-validation with 5 folds of the test and training datasets was conducted. Table 1 presents the MAE and CRPS values obtained across validation folds for all models.

Table 1. Average MAE and CRPS across validation folds. All metrics are normalized by lifetime length. The $\pm$ sign indicates the standard deviation of the metrics across folds.

	MAE	**CRPS**
Baseline HSMM	3.06 ± 1.83	1.96 ± 1.09
Ideal Q-HSMM	0.21 ± 0.03	0.27 ± 0.02
KNN Q-HSMM	3.11 ± 1.23	1.96 ± 0.92

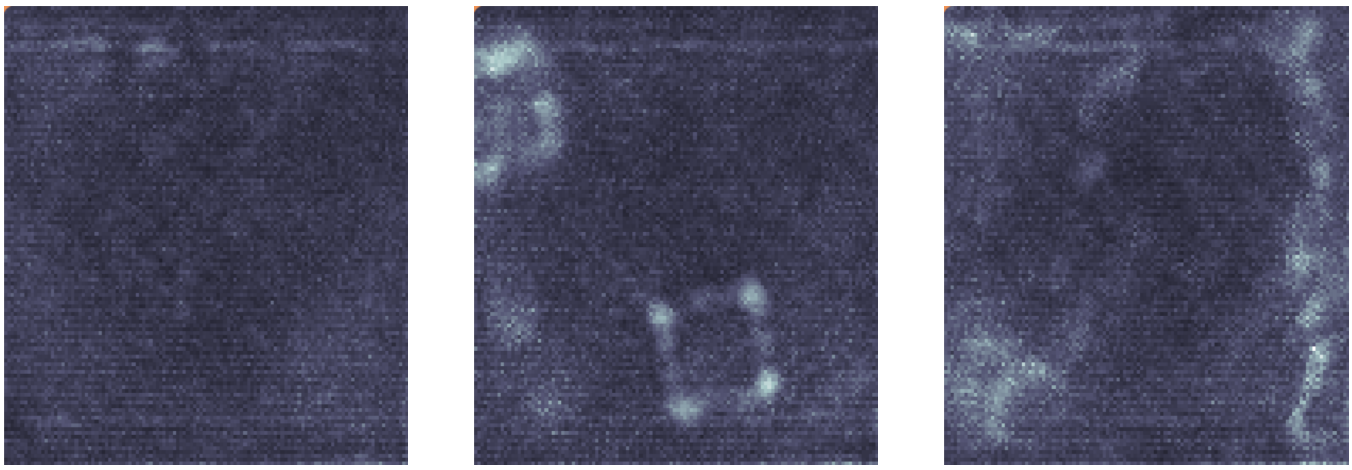

(a) Example of MQC data for a pristine specimen.

(b) Example of MQC data for a specimen with Teflon-induced defects.

(c) Example of MQC data for a specimen with oil-induced defects.

Figure 2. Examples of the employed MQC data for the MQ types available.

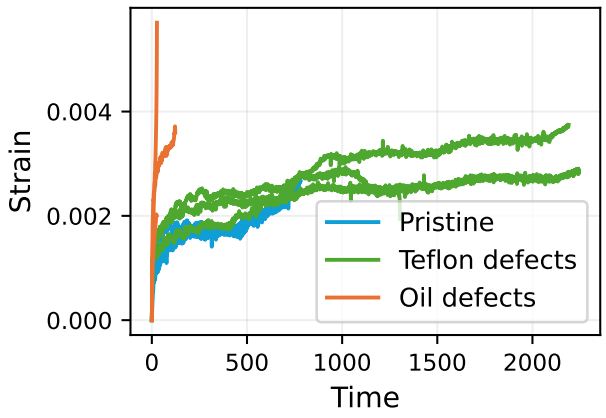

Figure 3. Raw strain data used for testing the framework. Each color corresponds to one MQ type.

Results in Table 1 show that managing past uncertainty with the ideal Q-HSMM leads to a drastic improvement in MAE and CRPS, when compared with the baseline HSMM. In fact, the MAE and CRPS show average improvements of 93% and 86%, respectively. This implies that the proper implementation of the framework can drastically reduce prognostic uncertainty while simultaneously reducing bias. These results are consistent across folds, as the metric standard deviations are negligible. Indeed, such a drastic reduction enables the prediction of reliable prognostics that lead to actionable decision-making.

However, the improvement observed in the ideal Q-HSMM does not carry over to the KNN Q-HSMM. This result can be attributed to the performance of the MQ type classifier. For this study, two further aspects are evaluated: classification accuracy (Figure 7) and certainty (Figure 8). Accuracy reflects the classification bias, while certainty indicates how much probability is assigned to the true MQ type label.

To evaluate classification accuracy, Figure 7 shows the confusion matrix for the KNN classifier. From the figure, it is evident that the classifier struggles to differentiate between pristine systems and those contaminated with oil defects. Indeed, the obtained classification accuracy is only 60%. This implies that, in most cases, the framework will assign greater weight to the wrong Q-HSMM, leading to bias in the prognostic.

In addition, Figure 8 illustrates how the KNN classifier assigns probabilities to the true MQ type label for all specimens across folds. The figure shows that, for most classifications, the assigned probabilities are equivalent to a random guess. This indicates that the classifier is not adding certainty to the prognostics and is thus failing at past uncertainty management.

These results support the claim that counting with an accurate MQ type classifier is crucial to effective past uncertainty management. In fact, if the MQ type classifier is mistaken, this error will propagate to the prognostic after performing past uncertainty management. Moreover, even if the classifier is accurate, it should also be confident in its predictions. Otherwise, it will not be able to add certainty in the prognostics.

5. CONCLUSIONS

This work proposes a framework for managing past uncertainty in prognostics. The framework uses MQC data to select Q-HSMMs, trying ot render out the contribution of past uncertainty from the output prognostics. The proposed framework constitutes an initial approach towards incorporating

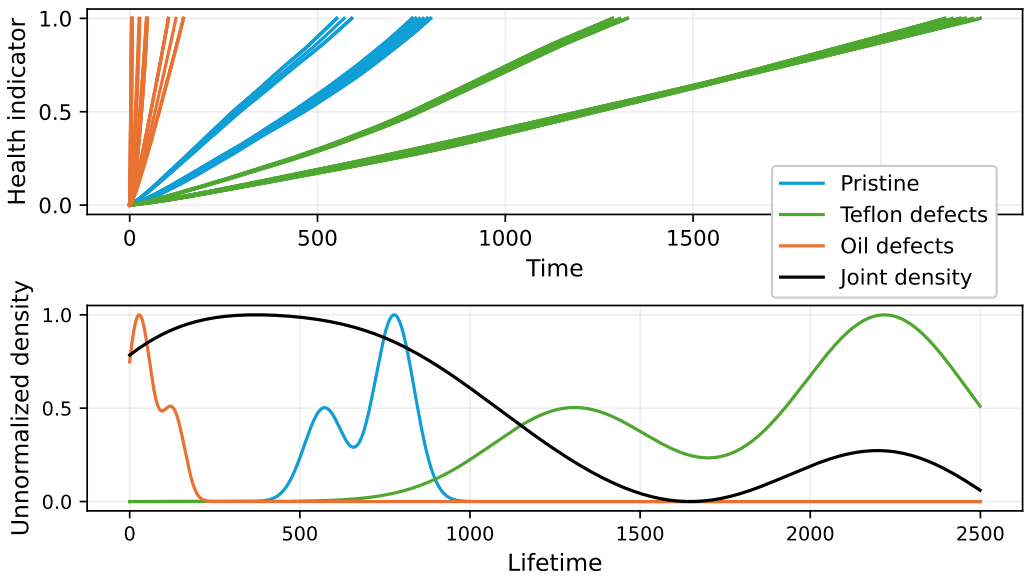

Figure 4. a) Employed degradation histories, separated by their corresponding defect type. b) Scaled density distributions for the specimen lifetime, separated by their corresponding defect type. The black density corresponds to the specimens lifetime, without distinguishing between manufacturing defects.

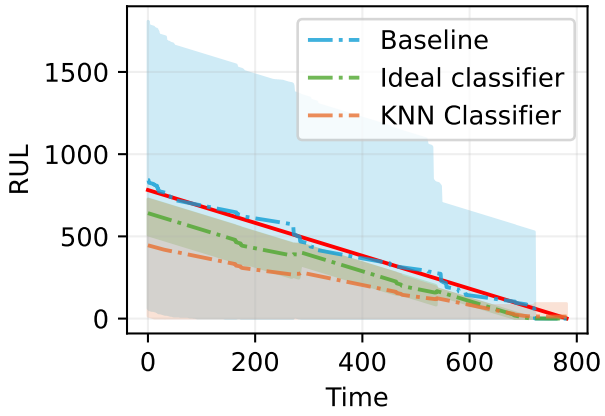

Figure 5. Actual RUL for the implemented prognostic models in a test pristine specimen.

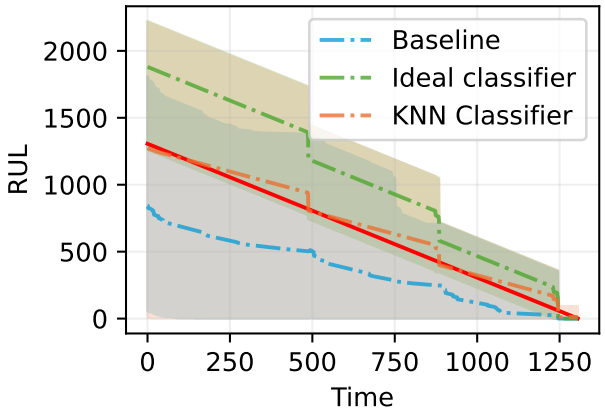

Figure 6. Actual RUL curve for the implemented prognostic models in a test specimen with Teflon-induced defects.

past uncertainty management in PHM solutions.

The framework was implemented using data from composite specimens with different MQ types cycled under fatigue. Two variants of the framework were tested: one uses an ideal MQ type classifier, while the other uses a KNN classifier. The performance of these variants is benchmarked against a baseline HSMM.

Results indicate that using the ideal classifier for past uncertainty management leads to a consistent reduction in uncertainty and an increase in overall prognostics accuracy compared to the baseline HSMM. This outcome showcases the potential benefits of incorporating past uncertainty management to enhance prognostics reliability. For the KNN classifier, the results show a drastic drop in performance compared to the ideal case. This drop in performance is attributed to the classifier, suggesting that accurate MQ type classification is fundamental for effective past uncertainty management.

Future work will focus on understanding which are the necessary requirements that guarantee effective past uncertainty management. These requirements could later be enforced in the HSMM and MQ type classifier, enhancing the robustness

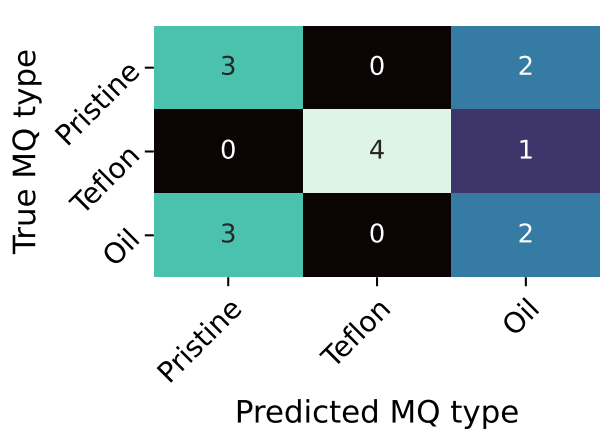

Figure 7. Confusion matrix for the KNN classifier.

of the framework. Additionally, future research will explore how to incorporate the proposed framework in cases where other sources of uncertainty are also managed, with the potential for a complete framework for UM in prognostics.

ACKNOWLDEGMENTS

The authors would like to acknowledge Ir. Dennis Leinarts for providing the dataset used as a case study.

This work has been funded by the IntelliWind Doctoral Network, granted under the Horizon Europe MSCA Doctoral Network programme, Grant agreement 101168725.

REFERENCES

Edwards, D., Orchard, M. E., Tang, L., Goebel, K., & Vachtsevanos, G. (2010). Impact of input uncertainty on failure prognostic algorithms: Extending the remaining useful life of nonlinear systems, 2010. In *Annual conference of the prognostics and health management society.*

Eleftheroglou, N. (2020). *Adaptive prognostics for remaining useful life of composite structures* (Dissertation (TU Delft), Delft University of Technology). doi: 10.4233/uuid:538558fb-ac9a-414d-8a59-4b523d8ff74c

Eleftheroglou, N., Galanopoulos, G., & Loutas, T. (2024, 3). Similarity learning hidden semi-markov model for adaptive prognostics of composite structures. *Reliability Engineering & System Safety*, *243*, 109808. doi: 10.1016/j.ress.2023.109808

Eleftheroglou, N., Zarouchas, D., & Benedictus, R. (2020, 8). An adaptive probabilistic data-driven methodology for prognosis of the fatigue life of composite structures. *Composite Structures*, *245*, 112386. doi: 10.1016/j.compstruct.2020.112386

Endrerud, J. O. (2014, 7). Ndt for non-experts. *Reinforced Plastics*, *58*, 44-46. doi: 10.1016/S0034-3617(14)70184-5

Goli, M., Ghodrati, B., & Eleftheroglou, N. (2025, 12). A literature review-based evaluation framework for maintenance strategy selection in heavy vehicles. *Results in Engineering*, *28*, 107109. doi: 10.1016/j.rineng.2025.107109

Kim, S., Choi, J.-H., & Kim, N. H. (2022, 6). Inspection schedule for prognostics with uncertainty management. *Reliability Engineering & System Safety*, *222*, 108391. doi: 10.1016/j.ress.2022.108391

Kontogiannis, T., Salinas-Camus, M., & Eleftheroglou, N. (2025). Hidden markov model applications. In *Stochastic modeling and statistical methods* (p. 191-213). Elsevier. doi: 10.1016/B978-0-44-331694-4.00015-3

Leinarts, D. R. (2025). *Uncertainty management for past state prognostic uncertainties in aerospace structures* (Master's thesis). TU Delft.

Orchard, M., Kacprzynski, G., Goebel, K., Saha, B., & Vachtsevanos, G. (2008, 10). Advances in uncertainty representation and management for particle filtering applied to prognostics. In *2008 international conference on prognostics and health management* (p. 1-6). IEEE. doi: 10.1109/PHM.2008.4711433

Saha, B., & Goebel, K. (2008, 3). Uncertainty management for diagnostics and prognostics of batteries using bayesian techniques. In *2008 ieee aerospace conference* (p. 1-8). IEEE. doi: 10.1109/AERO.2008.4526631

Salinas-Camus, M., & Eleftheroglou, N. (2024, 6). Uncertainty in aircraft turbofan engine prognostics on the c-mapss dataset. *PHM Society European Conference*, *8*, 10. doi: 10.36001/phme.2024.v8i1.4007

Salinas-Camus, M., Goebel, K., & Eleftheroglou, N. (2025, 8). 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

Sankararaman, S. (2015, 2). Significance, interpretation, and quantification of uncertainty in prognostics and remaining useful life prediction. *Mechanical Systems and Signal Processing*, *52-53*, 228-247. doi: 10.1016/j.ymssp.2014.05.029

Sankararaman, S., & Goebel, K. (2013, 10). Why is the remaining useful life prediction uncertain? *Annual Conference of the PHM Society*, *5*. doi: 10.36001/phmconf.2013.v5i1.2263

Syriopoulos, P. K., Kalampalikis, N. G., Kotsiantis, S. B., & Vrahatis, M. N. (2025, 2). knn classification: a review. *Annals of Mathematics and Artificial Intelligence*, *93*, 43-75. doi: 10.1007/s10472-023-09882-x

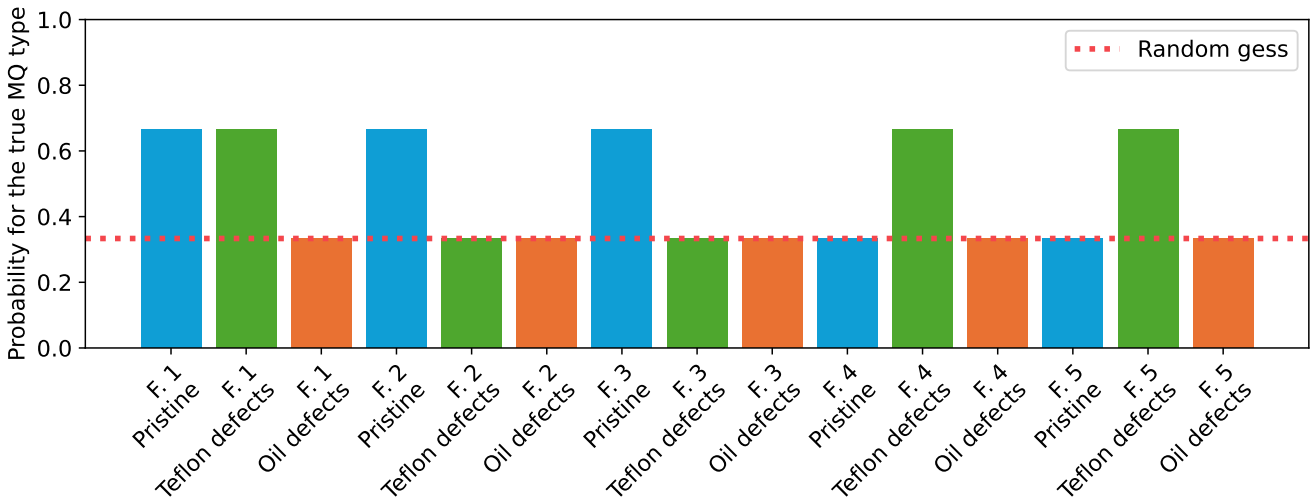

Figure 8. Probabilities assigned to the true MQ type for the KNN classifier across folds. The symbol 'F.' indicates the fold-number, followed by the MQ type of the tested system.

Tang, L., Kacprzynski, G. J., Goebel, K., & Vachtsevanos, G. (2009, 3). Methodologies for uncertainty management in prognostics. In *2009 ieee aerospace conference* (p. 1-12). IEEE. doi: 10.1109/AERO.2009.4839668