Proceedings of the 35th European Safety and Reliability & the 33rd Society for Risk Analysis Europe Conference
Edited by Eirik Bjorheim Abrahamsen, Terje Aven, Frederic Bouder, Roger Flage, Marja Ylönen

doi: 10.3850/978-981-94-3281-3_ESREL-SRA-E2025-P2692-cd

Decision-Focused Predictive Maintenance: Bridging the Gap between Data-Driven RUL Prognostics and Maintenance Planning

Kevin Wang
CentraleSupélec, Université Paris-Saclay, France
E-mail: kevin.wang@student-cs.fr

Zhuojun Xie
Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France
E-mail: xie.zhuojun@centralesupelec.fr

Mihaela Mitici
Faculty of Science, Utrecht University, The Netherlands
E-mail: m.a.mitici@uu.nl

Adam Abdin
Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France
E-mail: adam.abdin@centralesupelec.fr

Yiping Fang
Laboratoire Génie Industriel, CentraleSupélec, Université Paris-Saclay, France
E-mail: yiping.fang@centralesupelec.fr

Data-driven predictive maintenance (PdM) increasingly leverages machine learning techniques to predict remaining useful life (RUL) using abundant sensor data, supporting effective maintenance planning. However, most existing research follows a *predict-then-optimize* (PtO) paradigm, focusing on prognostic accuracy while overlooking how RUL predictions affect maintenance decisions. We propose a novel *Decision-Focused Predictive Maintenance* framework that bridges the gap between RUL prognostics and maintenance planning. This framework creates an end-to-end pipeline that directly connects RUL estimation to maintenance actions. An experiment using the CMAPSS dataset demonstrates that our framework achieves a 9.3% reduction in maintenance costs compared to the PtO approach. This improvement is primarily attributed to the avoidance of unnecessary preventive maintenance, leading to a reduction in average lifetime waste due to preventive maintenance from 20.9 to 11.3 cycles. More importantly, we highlight the distinction between DFPdM and PtO by analyzing the quantile levels of RUL labels and maintenance decisions, demonstrating that DFPdM exhibits greater consistency in unifying estimation and optimization. Interestingly, we also observe that DFPdM achieves an acceptable prognostic accuracy, despite not being the primary training objective. This prognostics accomplished by recalibrating a specific quantile of the estimated distribution, rather than relying on the expectation or median as is common in conventional approaches.

Keywords: Predictive Maintenance, Prognostics, Decision-Focused Learning, Contextual Decision-Making

1. Introduction

Maintenance optimization has gained significant attention in both academia and industry as a critical strategy for reducing operational and maintenance costs, particularly in response to emerging challenges in the era of big data (De Jonge and Scarf, 2020; Pinciroli et al., 2023). The integration of sensors into modern machinery and systems, combined with advances in machine learning (ML), has driven the rise of data-driven predictive maintenance (PdM) (Wen et al., 2022). Prognostics, central to this approach, focuses on predicting the remaining useful life (RUL) of machinery using data analytics with sensors.

Typically, data-driven predictive maintenance adopts a *Predict-then-Optimize* (PtO) paradigm. In PtO, machine learning models are pretrained to generate RUL prognostics, which subsequently serve as input to the downstream maintenance optimization. The decision-making model uses these prognostics as input to prescribe maintenance decisions based on a specified criterion. This methodology has shown significant effectiveness in reducing maintenance costs and minimizing failure rates in practical settings (Mitici et al., 2023). However, the PtO paradigm has been shown to produce suboptimal decisions when the model class is misspecified (Hu et al., 2022). This suboptimality stems from two key issues. First, predefined model classes often fail to cover the true data-generating process, leading to inherent biases. Second, these models are typically optimized for predictive accuracy rather than decision quality, resulting in a misalignment between prognostic objectives and decision-making needs (Kong et al., 2022). To further improve predictive maintenance by reducing costs and preventing failures, it is important to bridge the gap between RUL prognostics and maintenance scheduling (JDMD Editorial Office et al., 2023). This work seeks to fill the gap by developing a closed-loop, end-to-end framework inspired by recent advances in Decision-Focused Learning Mandi et al. (2024).

1.1. *Related Works*

Several studies use the PtO approach for predictive maintenance. de Pater et al. (2022) applied a Convolutional Neural Network (CNN) to predict point RUL, which was subsequently employed as input for an integer programming-based maintenance plan. Similarly, Chen et al. (2021) utilized an LSTM to predict point RUL and made maintenance decisions with a threshold-based policy. When multiple uncertainty sources arise, probabilistic models become more appealing. For instance, Lee and Mitici (2023) trained a CNN with Monte Carlo dropout to generate probabilistic RUL estimates, which informed a deep reinforcement learning policy for multi-stage maintenance planning. Further, Mitici et al. (2023) developed single and multiple component replacement models based on a renewal-reward process to determine maintenance actions once the RUL distribution was estimated by a CNN model.

In these studies, the predictive models were trained separately from maintenance planning using regression losses, such as Mean Squared Error (MSE), which are symmetric with respect to the prediction error. However, maintenance costs are typically asymmetric in terms of under- and over-estimation. In contrast, our framework integrates predictive modeling and decision-making, thus directly improves the maintenance effectiveness.

Another stream of research focuses on the assessment metrics for various prognostic algorithms (Lewis and Groth, 2022), which share a similar essence with our work. It is recognized that when evaluating the performance of different prognostic models, the impact on downstream decision-making and health management tasks should be considered (Atamuradov et al., 2017). However, these metrics are typically designed for model comparison after training rather than for model training. There are two key reasons for this. First, forecast accuracy has traditionally been the primary objective of prognostics. Second, these metrics often involve parametric optimization or policy, creating technical barriers to their use in model training. One exception is Kamariotis et al. (2024), which conducted a broad review on data-driven predictive maintenance and introduced a decision-oriented metric for evaluating prognostic algorithms. The metric was used to tune hyperparameters in the model training process. However, tuning these hyperparameters in this manner can be computationally expensive. In contrast, our framework directly incorporates a decision-oriented metric for training the model's parameters instead of hyperparameters.

Our contributions are as follows. First, we propose the first decision-focused predictive maintenance framework that seamlessly integrates predictive and decision-making models into a closed-loop pipeline prioritizing decision quality. Second, we instantiate our framework by an integration of a Heteroscedastic Neural Network (HNN) and a differentiable stochastic policy to mitigate the impacts of model uncertainty on decision-making, and the computational expenses in the optimization differentiation, respectively. Third, we compare our framework against the PtO approach on the CMAPSS dataset, showing our framework reduces maintenance costs by effectively avoiding unnecessary interventions with slight compromise on the prognostic accuracy.

2. Methodology

2.1. *Predictive Module*

Throughout this work, we assume the specific downstream task of RUL prognostics is to prescribe maintenance decisions. Additionally, we consider the maintenance actions belong to a discrete feasible set $[H] \coloneqq \{0, 1, \ldots, H-1\}$ denoting H unit-time intervals.

An offline dataset $\mathcal{D}_n \coloneqq (s_i, T_i)_{i=1}^n$ is collected, consisting of n pairs of sensor data s_i and their corresponding RUL labels T_i. Additionally, the survival time k_i may be recorded to indicate how much time the equipment has been in operation when the sensor data were collected. The dataset can be derived from run-to-failure experiments. A predictive model $m(\cdot;\theta)$, parametrized by θ, is selected to estimate the RUL distribution over the same support $[H]$ based on the sensor data, i.e.,

$$\hat{P}_i = m(s_i;\theta), i = 1, 2, \ldots, n. \tag{1}$$

The parameter H represents the number of time intervals. A larger H corresponds to a longer look-ahead horizon, though it may increase computational complexity. The estimated distribution $\hat{P}_i$ is represented as a vector:

$$\hat{P}_i \coloneqq [\hat{p}_i^0, \hat{p}_i^1, \ldots, \hat{p}_i^{H-1}], \tag{2}$$

$$\sum_{h=0}^{H-1} \hat{p}_i^h = 1, \hat{p}_i^h \geq 0 \tag{3}$$

with $\hat{p}_i^h$ the estimated probability of the equipment failing in the time interval $[h, h+1)$. Notably, the probabilistic model can naturally revert to a deterministic model by assigning a value of 1 to a single $\hat{p}_i^h$ for a unique $h \in [H]$. For a compact representation, we denote the probability simplex as $\Delta_{H-1} \coloneqq \{p | \sum_{h=0}^{H-1} p^h = 1, p^h \geq 0\}$.

Various probabilistic prognostic methods exist. However, extending the look-ahead horizon can lead to high-dimensional distributions, making direct probability estimation computationally expensive. To address this, we adopt the HNN for generating probabilistic prognostics (Kamariotis et al., 2024) for its simplicity and model capacity. Computationally, a HNN model $m(\cdot;\theta)$ takes the sensor data s_i as input and output two elements $\hat{\mu}_i$ and $\hat{\sigma}_i$, mean and standard deviation, for each sample, which are further leveraged to construct a truncated Gaussian distribution $\hat{P}_i$ over support $[H]$,

$$\tilde{P}_i^h = \frac{1}{\sqrt{2\pi}\hat{\sigma}_i} e^{-(h-\hat{\mu}_i)^2/2\hat{\sigma}_i^2}, \tag{4}$$

$$\hat{P}_i^h = \frac{\tilde{P}_i^h}{\sum_{h'=0}^{H-1} \tilde{P}_i^{h'}} \tag{5}$$

This method reduces the dimensionality of the output from H to 2. However, it may introduce biases by restricting the model family as Gaussian ones.

2.2. *Planning Module*

With the estimated distributions, the modeler may seek an optimization method to determine the optimal maintenance time. In this work, we instantiate our framework with a single component replacement problem given constant preventive and corrective maintenance costs c_p and c_c, respectively. The model corresponds to a setting that the decision maker may inspect the system at arbitrary time and must make maintenance planning in advance. Our planning model follows a standard contextual stochastic optimization formulation (Sadana et al., 2025), which can be written as:

$$\begin{aligned} \hat{z}^*(s_i) \coloneqq & \arg\min_{z\in[H]} \mathbb{E}_{Y\sim m(S;\theta)}[c_i(z;Y)|S = s_i] \\ = & \arg\min_{z\in[H]} \mathbb{E}_{Y\sim \hat{P}_i}[c_i(z;Y)] \end{aligned} \tag{6}$$

where the feasible set is defined as $[H]$ when no additional constraints are imposed, and Y represents the random variable for the RUL, which follows the estimated conditional distribution given covariate $S = s_i$. The function c denotes the scenario-specific objective function, defined as:

$$c(z;y) = \frac{c_p \mathbb{I}[z \leq y]}{z + k_i} + \frac{c_c \mathbb{I}[z > y]}{y + k_i}, \tag{7}$$

with y a realization of Y, and $\mathbb{I}[\cdot]$ the indicator function. This scenario-wise objective function captures the preventive and corrective maintenance scenarios when the decision z and RUL y are given. An expectation over $\hat{P}_i$ is taken to quantify the expected Maintenance Cost per Unit Time (MCUT). Notably, this approach maintains the linearity of the value functions with respect to the estimated distribution, facilitating the computational efficiency of our framework. Compared to renewal process-based models, the contextual stochastic model emphasizes the uniqueness of each sample as captured by sensor data. This distinction is crucial for sensor-driven maintenance planning since sensor observations are inherently one-time events. That is, the probability of observing two identical sensor data is nearly zero, making a sample-specific approach essential for accurate decision-making.

2.3. *Hard Top-1*

To facilitate our framework, we propose an equivalent formulation to Eq. (6), which is essentially a parametric linear program. First, we define the concept of a value function for each decision as:

$$\hat{V}_i^h := \mathbb{E}_{Y\sim\hat{P}_i}[c(z=h;Y)], \tag{8}$$

$$\hat{V}_i := [\hat{V}_i^0, \hat{V}_i^1, \ldots, \hat{V}_i^{H-1}]^\top, \tag{9}$$

with $\hat{V}_i^h$ the objective value by substituting decision $z=h$ given the estimated distribution $\hat{P}_i$ for instance i. With the value function, the optimal decision $\hat{z}^*(\hat{P}_i)$ can be redefined as follows:

$$\hat{z}^*(\hat{P}_i) = [H]^\top w_i^*, \tag{10}$$

$$w_i^* := \underset{w\in\Delta_{H-1}}{\arg\min}\, \hat{V}_i^\top w, \tag{11}$$

where Eq. (11) finds the dimension of $\hat{V}_i$ that has the minimal value. By the property of the linear program, w_i^* is guaranteed when only one entry is 1, while the others are 0. By the inner product of $[H]$ and w_i^*, the optimal decision time is retrieved. It is straightforward to see that $\hat{V}_i$ is linear in $\hat{P}_i$ by the property of the expectation operator. This linear program formulation facilitates our framework by providing the gradient $\nabla_{\hat{P}_i}\hat{z}_i^*$ during backpropagation, as described in Section 3.2.

3. Decision-Focused Framework

3.1. *Pipeline*

Existing studies on data-driven predictive maintenance often follow a PtO pipeline. For this, model is first trained using the following formulation:

$$\min_{\theta\in\Theta} \frac{1}{n}\sum_{i=1}^{n} \ell(m(s_i;\theta), T_i) + \lambda\|\theta\|_2, \tag{12}$$

where Θ denotes the space of model parameters, $\lambda \geq 0$ the strength of regularization, and ℓ a specific regression loss. Depending the type of model and labels, MSE and Negative Log-Likelihood (NLL) losses are frequently adopted. Once parameter θ^* is obtained, the model will provide RUL estimations for unseen sensor data.

Our framework modify the regression loss ℓ for prognostics $m(s_i;\theta)$ in the PtO paradigm with a decision-oriented loss, known as regret. The sample-wise regret L_i is defined as follows:

$$\begin{aligned} L_i(\hat{P}_i, T_i) :=& \frac{c_p\mathbb{I}[\hat{z}_i^*(\hat{P}_i) \leq T_i]}{\hat{z}_i^*(\hat{P}_i) + k_i} \\ &+ \frac{c_c\mathbb{I}[\hat{z}_i^*(\hat{P}_i) > T_i]}{T_i + k_i} \\ &- \frac{c_p}{T_i + k_i}, \end{aligned} \tag{13}$$

which directly evaluates the decision quality derived by estimation $\hat{P}_i$ under the objective function c_i with RUL realization T_i, while the perfect MCUT is regarded as the baseline. Consequently, our framework aims to solve the following regularized empirical regret minimization problem by tuning the predictive model:

$$\min_{\theta\in\Theta} \frac{1}{n}\sum_{i=1}^{n} L_i(m(s_i;\theta), T_i) + \lambda\|\theta\|_2, \tag{14}$$

The structural difference between our framework and the PtO paradigm is illustrated in Fig. 1.

3.2. *Differentiation*

For data-driven predictive maintenance, the predictive models usually considered are neural networks since the sensor data can be temporal-dependent and high-dimensional. To provide an informative gradient during training with the

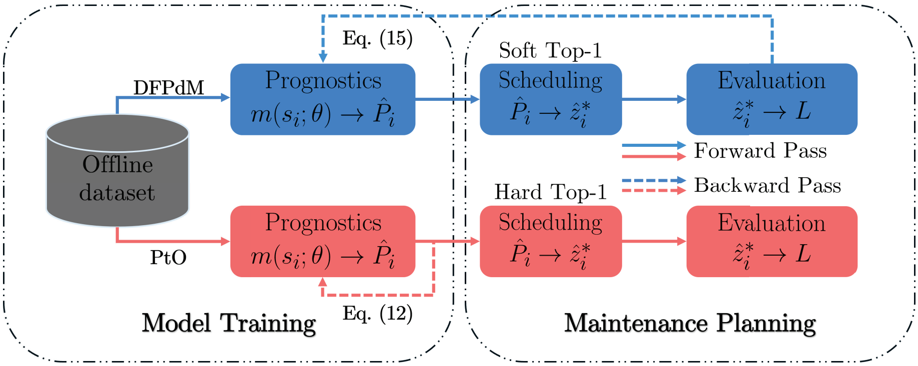

Fig. 1. Structural difference between Decision-Focused Predictive Maintenance (DFPdM) and PtO.

decision-oriented loss in Eq. (14) for back-propagation, we apply the following chain rule:

$$\nabla_\theta L = \frac{\partial L}{\partial \hat{z}^*}\frac{\partial \hat{z}^*}{\partial w^*}\frac{\partial w^*}{\partial \hat{V}}\frac{\partial \hat{V}}{\partial \hat{P}}\frac{\partial \hat{P}}{\partial \theta}, \tag{15}$$

in which $\frac{\partial \hat{z}^*}{\partial w^*} = [H]$, $\frac{\partial \hat{P}}{\partial \theta}$ is computed by auto-differentiation, and $\frac{\partial \hat{V}}{\partial \hat{P}}$ can be easily computed according to Eq. (8). The gradient of the regret with respect to the decision $\hat{z}^*$ and the gradient of weights w^* with respect to the value function $\hat{V}$ are challenging to obtain, which are addressed by methods proposed in the next section.

3.2.1. *Differentiable Regret*

For term $\frac{\partial L}{\partial \hat{z}^*}$, the gradient is constantly zero once decision $\hat{z}^*$ is larger than the perfect action time T_i. To avoid the gradient vanishing while considering the risk-aversion against corrective maintenance, we introduce the following surrogate regret loss for L_i:

$$\hat{L}_i(z,T) = \begin{cases} \frac{c_p}{z+k_i} - \frac{c_p}{T_i+k_i}, & \text{if } z \le T, \\ \frac{\alpha c_p}{2T_i - z + k_i} - \frac{c_p}{T+k_i}, & \text{otherwise}, \end{cases} \tag{16}$$

with $\alpha > 1$ a parameter that controls the aversion against corrective maintenance during training.

3.2.2. *Soft Top-1*

We consider a differentiable optimization layers for computing the term $\frac{\partial w^*}{\partial \hat{V}}$. Specifically, we consider the following soft Top-1 problem:

$$\tilde{w}^* = \arg\min_{w \in \Delta_{H-1}} \hat{V}^\top w + \epsilon R(w), \tag{17}$$

where $\epsilon R(w)$ is a convex regularization in w that avoids w from being 1. We take the quadratic regularization proposed by Amos and Kolter (2017) for computational tractability, i.e., $R(w) = w^\top w$, and implement it with package *Cvxpylayer* (Agrawal et al., 2019). Note that with this regularization, the values of the decision $\tilde{w}^*$ are no longer 0 or 1, but fractional numbers. Thus, the soft Top-1 policy becomes a surrogate to Eq. (10). The surrogate decision is determined as follows:

$$\tilde{z}^* = \lfloor [H]^\top \tilde{w}^* \rfloor, \tag{18}$$

which is a truncated integer by the floor operation $\lfloor \cdot \rfloor$ to mitigate the impact of regularization.

4. Experiments

We conduct a numerical experiment based on the subset **FD001** of the CMAPSS dataset. We follow the previous work Mitici et al. (2023) and choose the same number of sensors and length of observation window. We take the first 80 out of 100 engines from the **train_DF001.txt** file as training while the remaining 20 are used for testing. Moreover, we focus on the data with RUL less than $T_{th} = 100$ to underscore the importance of maintenance planning at later stages of equipment's lifetime. We adopt a linear neural network with two hidden layers as the prognostic model. All experiment details are summarized in Table 1.

We select the parameters $c_c = 500$ and $c_p = 100$ in the maintenance optimization for both PtO and DFPdM frameworks. Note that these parameters serve the framework in the policy. The con-

Table 1. Experimental details.

Framework	DFPdM	PtO
H	125	125
T_{th}	100	100
Hidden layer	[400,100]	[400,100]
Optimizer	Adam	Adam
Batchsize	32	32
Epochs	1500	1500
Learning rate	1e-4	1e-4
ϵ	0.05	-
α	5	-

figuration is thus fixed in forward pass because DFPdM integrates the predictive model and the policy. In the evaluation step, the decision quality is assessed by the average relative MCUT over the testing dataset $\mathcal{D}_{test}$:

$$M = \frac{1}{|\mathcal{D}_{test}|} \sum_{(s_i,T_i)\in\mathcal{D}_{test}} \frac{L_i}{c_p/(T_i + k_i)}. \quad (19)$$

Fig. 2 presents the prognostics and maintenance decisions computed by the two frameworks for five testing engines, indexed from 81 to 85. The MSE-PtO approach provides a reliable estimate of the remaining useful life (RUL) when using the quantile $q_{0.5}$ as the prognostic indicator, resulting in a mean absolute error (MAE) of 7.30 over the testing set. However, the stochastic policy derived from the MSE-PtO prognostics yields conservative decisions, leading to an MAE of 20.9 when compared to the ground truth RUL labels.

In contrast, the DFPdM framework prescribes less conservative decisions than PtO, though it remains conservative compared to the labels, especially at the beginning of the horizon. Overall, the conservativeness of DFPdM decisions can be attributed to the asymmetry of the regret function $\hat{L}$. During training, DFPdM informs the predictive model of the consequences associated with the prognostics. As a result, the MAE between DFPdM decisions and ground truth is reduced to 11.3. Notably, unlike PtO, where the median quantile $q_{0.5}$ is used as the prognostic, the optimal prognostic for DFPdM is selected as $q_{0.05}$. This quantile $q_{0.05}$ is first recalibrated on the training set before being applied to the testing set. As a byproduct of DFPdM, the MAE of DFPdM prognostics is 7.78. While slightly compromising prognostic accuracy, DFPdM reduces the MCUT by 9.3% compared to PtO. Moreover, it results in 12 corrective maintenance actions out of 2020 samples, compared to 5 made by PtO. It is important to note that the substantial reduction in RUL waste—from 20.9 to 11.3—only moderately impacts cost reduction, as the MCUT metric does not account for additional factors influencing preventive maintenance costs.

To comprehensively differentiate the two frameworks, we illustrate the distributions of quantile levels corresponding to ground truth and decisions in Fig. 3. As shown in Fig. (3a), both PtO and DFPdM decisions result in relatively consistent quantiles concentrated within the interval $[q_0, q_{0.02}]$, as both frameworks rely on the same stochastic optimization model. However, compared to PtO, DFPdM exhibits a narrower range of decision quantile levels, predominantly falling within $[q_{0.01}, q_{0.02}]$ for more than 85% of the samples. This highlights its consistency in decision-making and distribution estimation.

Furthermore, Fig. (3b) reveals the intrinsic differences between PtO and DFPdM in terms of prognostic estimation. By mapping RUL labels onto quantile levels, PtO demonstrates a nearly symmetric and even distribution around $q_{0.5}$, aligning with the calibration condition required for probabilistic models to effectively quantify uncertainty. In contrast, DFPdM adjusts the estimated distributions such that decisions are closer to the actual labels. This is driven by the principle that maintenance costs can be minimized when decisions approach the labels on the left side of the distribution. This observation also justifies the selection of $q_{0.05}$ as the optimal prognostic for the DFPdM framework, as it approximately partitions the samples into two equal subsets.

5. Conclusions

We proposed the first decision-focused framework for data-driven predictive maintenance incorporating ML models. The framework directly adopts a decision-oriented metric, i.e., regret, as the training objective for the ML-based predictive model

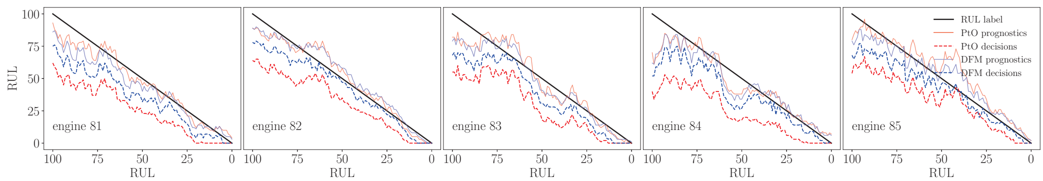

Fig. 2. Prognostics and decisions made by DFPdM and PtO for 5 out of 20 testing engines within 100 RUL. Note that PtO and DFM (short for DFPdM) take quantile $q_{0.5}$ and $q_{0.05}$, respectively.

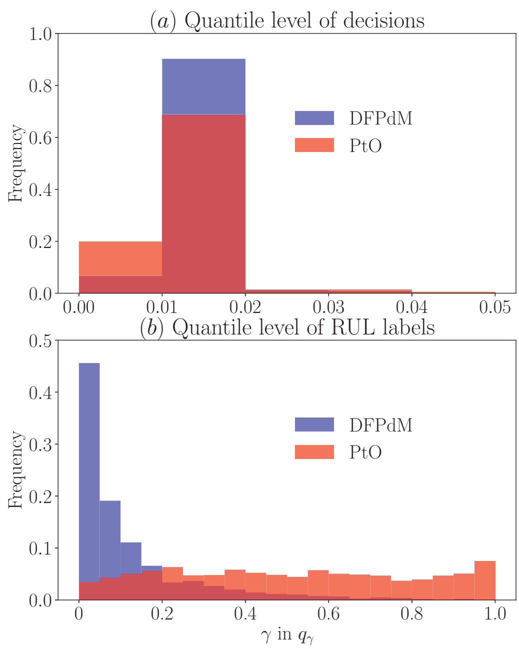

Fig. 3. Distribution of quantile levels in the testing samples for (a) RUL labels and (b) decisions from the two frameworks.

in contrast to PtO approaches which leverage a conventional regression loss. We instantiate our framework by seamlessly integrating probabilistic prognostics and a stochastic optimizer into a closed-loop, end-to-end pipeline. Experiments based on the CMAPSS dataset demonstrate the issue of target misalignment caused by PtO approaches. Instead, our framework generates interpretable prognostics and relatively conservative decisions. As such, our framework improves the average MCUT by 9.3%, reducing the average wasted RUL from 20.9 cycles to 11.3 cycles for 2020 samples. By analyzing the relationship between the quantiles, RUL labels and maintenance decisions, we demonstrate the consistency of DFPdM in prognostics and optimization and suggest that the traditional regression losses may be inappropriate for probabilistic prognostics if an explicit maintenance task follows.

Our framework is partially constrained by the capacity of the predictive model, which employs a two-layer neural network and heteroscedastic model and aims to capture the complex time-series characteristics of sensor data. Therefore, we limited ourself to implement the framework for engines with an actual RUL less than 100 cycles. We achieve improvements in the decision-oriented metric, as well as an increase in the computational cost to obtain maintenance decisions. This is primarily due to the need for solving multiple optimization instances during each forward pass and differentiating the soft Top-1 operator. Furthermore, as a preliminary work to demonstrate the effect of framework, our maintenance planning model considers limited factors. In the future works, we aim to address the computational challenges using more advanced ML models and conduct a more thorough analysis regarding the inconsistency issue of PtO approaches for predictive maintenance, as well as case studies incorporating realistic maintenance planning settings.

References

Agrawal, A., B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter (2019). Differentiable convex optimization layers. *Advances in neural information processing systems 32*, 9562–9574.

Amos, B. and J. Z. Kolter (2017). Optnet: Differentiable optimization as a layer in neural networks. In *International conference on machine learning*, pp. 136–145. PMLR.

Atamuradov, V., K. Medjaher, P. Dersin, B. Lamoureux, and N. Zerhouni (2017). Prognostics and health management for maintenance practitioners-review, implementation and tools evaluation. *International Journal of Prognostics and Health Management 8*(3), 1–31.

Chen, C., Z. H. Zhu, J. Shi, N. Lu, and B. Jiang (2021). Dynamic predictive maintenance scheduling using deep learning ensemble for system health prognostics. *IEEE Sensors Journal 21*(23), 26878–26891.

De Jonge, B. and P. A. Scarf (2020). A review on maintenance optimization. *European journal of operational research 285*(3), 805–824.

de Pater, I., A. Reijns, and M. Mitici (2022). Alarm-based predictive maintenance scheduling for aircraft engines with imperfect remaining useful life prognostics. *Reliability Engineering & System Safety 221*, 108341.

Hu, Y., N. Kallus, and X. Mao (2022). Fast rates for contextual linear optimization. *Management Science 68*(6), 4236–4245.

JDMD Editorial Office, N. Gebraeel, Y. Lei, N. Li, X. Si, and E. Zio (2023). Prognostics and remaining useful life prediction of machinery: Advances, opportunities and challenges. *Journal of Dynamics, Monitoring and Diagnostics 2*(1), 1–12.

Kamariotis, A., K. Tatsis, E. Chatzi, K. Goebel, and D. Straub (2024). A metric for assessing and optimizing data-driven prognostic algorithms for predictive maintenance. *Reliability Engineering & System Safety 242*, 109723.

Kong, L., J. Cui, Y. Zhuang, R. Feng, B. A. Prakash, and C. Zhang (2022). End-to-end stochastic optimization with energy-based model. *Advances in Neural Information Processing Systems 35*, 11341–11354.

Lee, J. and M. Mitici (2023). Deep reinforcement learning for predictive aircraft maintenance using probabilistic remaining-useful-life prognostics. *Reliability Engineering & System Safety 230*, 108908.

Lewis, A. D. and K. M. Groth (2022). Metrics for evaluating the performance of complex engineering system health monitoring models. *Reliability Engineering & System Safety 223*, 108473.

Mandi, J., J. Kotary, S. Berden, M. Mulamba, V. Bucarey, T. Guns, and F. Fioretto (2024). Decision-focused learning: Foundations, state of the art, benchmark and future opportunities. *Journal of Artificial Intelligence Research 80*, 1623–1701.

Mitici, M., I. de Pater, A. Barros, and Z. Zeng (2023). Dynamic predictive maintenance for multiple components using data-driven probabilistic rul prognostics: The case of turbofan engines. *Reliability Engineering & System Safety 234*, 109199.

Pinciroli, L., P. Baraldi, and E. Zio (2023). Maintenance optimization in industry 4.0. *Reliability Engineering & System Safety 234*, 109204.

Sadana, U., A. Chenreddy, E. Delage, A. Forel, E. Frejinger, and T. Vidal (2025). A survey of contextual optimization methods for decision-making under uncertainty. *European Journal of Operational Research 320*(2), 271–289.

Wen, Y., M. Fashiar Rahman, H. Xu, and T.-L. B. Tseng (2022). Recent advances and trends of predictive maintenance from data-driven machine prognostics perspective. *Measurement 187*, 110276.