

Article

A Comparative Study of Optimization Models for Condition-Based Maintenance Scheduling of an Aircraft Fleet

Iordanis Tseremoglou ^{1,*} , Paul J. van Kessel ²  and Bruno F. Santos ¹ ¹ Air Transport and Operations, Faculty of Aerospace Engineering, Delft University of Technology, 2629 HS Delft, The Netherlands² KLM Royal Dutch Airlines, 1117 ZL Schiphol, The Netherlands

* Correspondence: i.tseremoglou@tudelft.nl

Abstract: Condition-based maintenance (CBM) scheduling of an aircraft fleet in a disruptive environment while considering health prognostics for a set of systems is a very complex combinatorial problem, which is becoming more challenging in light of the uncertainty included in health prognostics. This type of problem falls under the broad category of resource-constrained scheduling problems under uncertainty and is often solved using a mixed integer linear programming (MILP) formulation. While a MILP framework is very promising, the problem size can scale exponentially with the number of considered aircraft and considered tasks, leading to significantly high computational costs. The most recent advances in artificial intelligence have demonstrated the capability of deep reinforcement learning (DRL) algorithms to alleviate this curse of dimensionality, as once the DRL agent is trained, it can achieve real-time optimization of the maintenance schedule. However, there is no guarantee of optimality. These comparative merits of a MILP and a DRL formulation for the aircraft fleet maintenance scheduling problem have not been discussed in the literature. This study is a response to this research gap. We conduct a comparison of a MILP and a DRL scheduling model, which are used to derive the optimal maintenance schedule for various maintenance scenarios for aircraft fleets of different sizes in a disruptive environment, while considering health prognostics and the available resources for the execution of each task. The quality of solutions is evaluated on the basis of four planning objectives, defined according to real airline practice. The results show that the DRL approach achieves better results with respect to scheduling of prognostics-driven tasks and requires less computational time, whereas the MILP model produces more stable maintenance schedules and induces less maintenance ground time. Overall, the comparison provides valuable insights for the integration of health prognostics in airline maintenance practice.

Keywords: condition-based maintenance (CBM); partially observable markov decision process, (POMDP); partially observable Monte Carlo planning (POMCP); deep reinforcement learning (DRL); mixed integer linear programming (MILP); planning under uncertainty



Citation: Tseremoglou, I.; van Kessel, P.J.; Santos, B.F. A Comparative Study of Optimization Models for Condition-Based Maintenance Scheduling of an Aircraft Fleet. *Aerospace* **2023**, *10*, 120. <https://doi.org/10.3390/aerospace10020120>

Academic Editor: Michael Schultz

Received: 31 October 2022

Revised: 16 January 2023

Accepted: 24 January 2023

Published: 27 January 2023



Copyright: © 2023 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Aircraft maintenance is one of the key contributors to aviation safety and operational efficiency in the airline industry. Total maintenance expenditure accounts for approximately 10–15% of the direct operating costs of an aircraft operator [1]. Traditional aircraft maintenance policies are based on a combination of the *preventive* and *corrective* approach.

According to the preventive strategy, components need to be inspected/replaced at fixed intervals, which come in the form of flight hours (FHs), flight cycles (FCs), or calendar days (DYs). This strategy is implemented through the preventive maintenance tasks described in the Maintenance Planning Document (MPD) and included in scheduled maintenance checks, referred to as letter checks (e.g., A-, B-, C-, and D-checks). While this strategy is largely responsible for the safety and reliability of aviation today, its cost-effectiveness is limited by statistical considerations and generalizations, that can lead to a

replacement of a component long before its true due date is reached, or to a failure of the component prior to the assigned maintenance date. In both cases, improved operational costs are induced.

The corrective maintenance strategy addresses unexpected maintenance tasks, such as faults reported by the pilots, bird strikes, or a finding during an inspection or a functional check as part of a preventive maintenance task. The stochastic nature of these unanticipated maintenance tasks creates disruptions to the maintenance schedule. As such, the maintenance schedule cannot be executed as planned but, rather, has to be constantly adjusted. In the best-case scenario, the aircraft can be assigned to a new maintenance opportunity. However, in the worst-case scenario, there might not be an available maintenance opportunity and the aircraft will lose its airworthiness until the issue is resolved. What is more, usually the resources required to repair the component after failure can be more costly than repairing the component before it fails.

Striving to reduce the related maintenance costs, the aviation industry is gradually shifting to condition-based maintenance (CBM). CBM is a predictive maintenance strategy that makes use of sensors and advanced analytics to continuously monitor the health of aircraft components and predict their remaining useful life (RUL). Maintenance actions are triggered only when there is strong evidence of failure risk, hence decreasing the number of unnecessary maintenance actions and, at the same time, avoiding unforeseen failures and corresponding unscheduled maintenance events. Since preventive maintenance strategies will remain necessary for an effective transition to CBM, this gradual transition is currently realized by combining the existing corrective and preventive with the predictive maintenance strategy.

However, this “hybrid” maintenance approach is very challenging for the maintenance planner. The first challenge arises from the uncertainty included in the RUL predictions. More specifically, the maintenance planner, instead of planning tasks considering a deterministic task due date, as in the case of the preventive and corrective maintenance tasks, has to devise the maintenance plan based on an uncertain outcome. This outcome is captured by the probability distribution of the predicted RUL.

The second challenge lies in the continuous update of the (uncertain) RUL predictions together with the stochastic arrival of unanticipated maintenance tasks, which eventually can create disruptions to the existing maintenance schedule. These disruptions compromise the feasibility and efficiency of the initial schedule. For this reason, the maintenance schedule needs to be continuously updated with respect to the new information, such that all maintenance tasks are scheduled ahead of their due date. Currently, the majority of airlines still rely on manual scheduling methods to adjust and update their maintenance schedule, which results, potentially, in sub-optimal planning.

Historically, the academic literature on aircraft maintenance scheduling is mostly oriented around the scheduling of maintenance tasks bundled in letter checks. Among the most recent works, ref. [2] used a Dynamic Programming (DP)-based methodology to solve the aircraft maintenance check scheduling problem over a period of 4 years. However, this level of scheduling is too abstract to address the disruptions mentioned above. To achieve the level of detail required for dynamic scheduling of maintenance tasks, some researchers have shifted their attention to scheduling maintenance tasks individually. The most recent and detailed work was performed by [3], who proposed a hybrid value function approximation (VFA)–Rolling Horizon policy to schedule critical and non-critical tasks of an aircraft fleet. However, the availability of resources, such as material or ground equipment requirements, is not considered.

In general, task (re-)scheduling in a disruptive environment while taking into account the availability of resources touches upon three fields of research: task scheduling, disruption management, and resource allocation. A widely used method in the literature for addressing these types of problems is mixed integer linear programming (MILP) algorithms ([4–6]). Although the literature on MILP frameworks specifically applied to the aircraft maintenance-scheduling problem is extremely scarce, there are other fields of

study where the developed models can be extended to the airline maintenance scheduling problem. Specifically, both the health and the construction sector show similarities to airline maintenance operations when it comes to scheduling and managing disruptions.

In the research area of the health sector, Ref. [7] developed a MILP model to address the issue of emergency patient admission and rescheduling of elective patients, subject to the availability of various resources, such as the capacity of clinic units. The objective of the MILP model is to minimize the cost of postponing the elective surgery patients, declining the emergency patients and overutilization of operating rooms. Ref. [8] used a MILP model to minimize the tardiness of surgeries, idle time, and overtime, while also considering uncertainty in the duration of the surgeries and constraints related to human resources and medical equipment. Ref. [9] design a MILP model that provides support for the selection and rescheduling (when needed) for elective patients waiting for surgery. The results of the model were evaluated with respect to four different scheduling objectives. In the construction sector, Ref. [10] proposed a rescheduling optimization model which minimizes alterations to the initial project schedule. The only contribution we are aware of, that addresses aircraft maintenance task scheduling in a disruptive environment, is [11], who developed a MILP framework to schedule aircraft tasks, minimizing aircraft ground time while also limiting the number of schedule changes.

The main issue exhibited in the previous studies is that the problem size scales exponentially with the number of the considered tasks and the related decision variables, which could take significant computational time and memory usage for applications in real-time environment. Reinforcement learning is able to alleviate this time limitation, as the time required to find a solution can be heavily reduced once the RL model is trained. The use of RL in aircraft maintenance scheduling problems was studied in [12], who develop a deep Q-learning network (DQN) to optimize the long-term scheduling maintenance of an aircraft fleet.

Most recently, in [13], a two-stage scheduling framework for an aircraft fleet in a CBM context is developed, taking into account the uncertainty of the RUL predictions of the prognostics-driven tasks. In the first stage, the proposed framework uses the partially observable Monte Carlo planning (POMCP) algorithm developed by [14], to define the optimal maintenance action for prognostics-driven tasks with uncertain RUL predictions. In the second stage, a deep reinforcement learning (DRL) algorithm is developed, which produces the maintenance schedule for the aircraft fleet, where each aircraft contains a mixture of preventive, corrective and prognostics-driven tasks. The obtained results indicate that DRL is able to produce both an efficient and stable maintenance schedule in just a few seconds.

In this paper, we adopt the framework developed in [13] and we further use in the second stage of scheduling the MILP model developed by [11], in order to compare and explore the performance and the possible trade-offs between the MILP and DRL scheduling approach.

The remainder of this paper is organized as follows: the aircraft maintenance requirements, objectives, and the general problem formulation are described in Section 2. The maintenance scheduling algorithm for the prognostics-driven tasks is briefly discussed in Section 3. An overview of the MILP and the DRL scheduling model for the aircraft fleet is described in Section 4. In Section 5, the performance of both models is evaluated on three different maintenance scenarios for different aircraft fleet sizes, where each aircraft has a list of open maintenance tasks that are updated on a continuous basis. Finally, Section 6 summarizes the research with concluding remarks and recommendations for future work.

2. Problem Formulation

2.1. Problem Definition and Scope

The problem we are addressing can be summarized as follows: let us consider an aircraft fleet R , where each aircraft $r \in R$, is having a set of monitored components, U_r . The maintenance of each monitored component, $u_r \in U_r$, $r \in R$, is driven by the probabilis-

tic RUL predictions of the prognostics model through the corresponding prognostics-driven tasks, G_{progn}^r . In total, the maintenance tasks G_r of each aircraft $r \in R$ consist of a subset of corrective maintenance tasks, G_{corr}^r , preventive maintenance tasks, G_{prev}^r , and prognostics-driven tasks, G_{progn}^r . For ease of reference, we provide a list of the frequently used notation in Nomenclature at the end of the paper.

This list of open maintenance tasks of each aircraft is continuously updated, either due to the unexpected arrival of corrective maintenance tasks or due to new RUL predictions generated by the prognostics model. Following this, the maintenance schedule has to be continuously adjusted to the new information, while, at the same time, the following four objectives need to be satisfied:

1. All tasks should be performed ahead of their due date,
2. Ground time required for maintenance should be minimized,
3. The number of rescheduling actions should be minimized,
4. Tasks should be scheduled at their optimal moment in time.

A more detailed explanation of the planning objectives is provided in Section 2.4. An overview of the modelling setup is presented in Figure 1. The modelling blocks are represented with a grey background. In order to translate the RUL predictions in prognostics-driven tasks with associated recommended maintenance dates based on the minimization of maintenance costs, we use the partially observable Monte Carlo planning algorithm as developed in [13]. Two different scheduling optimization models are used and compared, one based on the MILP framework developed by [11] and one utilizing the DRL algorithm developed in [13]. The modelling blocks are explained in more detail in the next sections.

The general concept is that the POMCP algorithm will consider the probabilistic RUL predictions generated from the RUL prognostics model to create belief states on the health degradation of the corresponding component. The belief state is a probability measure to estimate the state of the system. Taking into account these belief states, the available maintenance slots in the foreseen future, and the preventive and corrective maintenance costs for the considered components, the POMCP prescribes the recommended dates for the execution of the prognostics-driven tasks. These tasks, together with the list of tasks included in the Maintenance Planning Document (MPD) and the Minimum Equipment List (MEL), are added to the pool of tasks that the optimization models will consider. The schedule optimization models will assign these tasks to the available maintenance slots in a time horizon of multiple weeks, considering the available resources.




Figure 1. Modelling framework overview.

2.2. Inputs of the Scheduling Frameworks

A brief explanation of the inputs that are considered by the scheduling frameworks is provided below. Readers are referred to [13] for a detailed explanation of the inputs.

- **Predicted RUL distributions:** The distributions that capture the predicted amount of working hours left until the end-of-life of the component.
- **Maintenance slots:** The available maintenance opportunities, M , for execution of aircraft maintenance. These are further divided in *fixed* maintenance slots, M_{fixed} , which are tied to a specific aircraft registration, and *flexible* maintenance slots, $M_{flexible}$, which are tied to a specific aircraft type. Contrary to the flexible slots, the fixed slots are determined several weeks ahead in order to accommodate more extensive maintenance checks, such as the letter checks.
- **Maintenance costs:** These involve the preventive, C_{ur}^{prev} , and corrective maintenance costs, C_{ur}^{corr} , induced when executing a prognostics-driven task.
- **Resources:** The required *Material*, *Machinery*, *Method* and *Manpower* (4M) for the execution of each task. *Material* refers to the required consumable parts/spares required for the execution of the maintenance task, whereas *Machinery* corresponds to the required ground equipment. *Method* refers to the required ground time for the completion of the task, and *Manpower* captures the requirements in workhours needed for the execution of the task.
- **Current maintenance schedule:** The schedule at the present time point, that details the allocation of aircraft and related tasks to maintenance slots.
- **Prognostics-driven tasks:** The tasks that correspond to the maintenance of components that are monitored permanently through sensors and for which there is a predicted RUL distribution obtained from the prognostics model on a continuous basis.
- **Preventive and corrective tasks:** Preventive tasks are included in the Maintenance Planning Document (MPD) and are performed in fixed periodic inspection intervals that come in the form of FHs, FCs, or DYs. Corrective tasks are unexpected maintenance tasks, such as findings during the execution of a preventive maintenance task or a fault reported by the pilots. They have to be restored within a specific time window, which varies from a few days to few months.

2.3. Constraints

The execution of each maintenance task is constrained by the availability of maintenance resources in maintenance slot m . More specifically, a maintenance task cannot be executed before the required materials and equipment are available (*Material* and *Machinery*). This is ensured through the following set of constraints:

$$\sum_{m \in M} (1 - \text{Material}_{g,m}) \cdot T_{g,m} = 0 \quad \forall g \in G \quad (1)$$

$$\sum_{m \in M} (1 - \text{Machinery}_{g,m}) \cdot T_{g,m} = 0 \quad \forall g \in G \quad (2)$$

where $\text{Material}_{g,m}$ and $\text{Machinery}_{g,m}$ take unitary values if the required material and equipment for the execution of maintenance task g are available before the start date of maintenance slot m , and $T_{g,m}$ takes unitary value if task g is allocated to maintenance slot m .

Furthermore, the required time, Duration_g , and workhours of skill w , GR_g^w , for the execution of the task g , must not exceed the duration, Duration_m , and the available workhours of skill w , GR_m^w , of maintenance slot m . The former requirements correspond to *Method* and *Manpower* and are implemented through the following set of constraints:

$$\text{Duration}_g \cdot T_{g,m} \leq \text{Duration}_m \quad \forall m \in M, g \in G \quad (3)$$

$$\sum_{g \in G} GR_g^w \cdot T_{g,m} \leq GR_m^w \quad \forall w \in W, m \in M \quad (4)$$

2.4. Maintenance Planning Objectives

We define the maintenance planning objectives based on the mindset of a maintenance planner operating in a real airline environment. The considered objectives are described as follows:

1. **Timely task execution:** The first and most important objective is to execute the open tasks before their due date, as when this date is exceeded the aircraft loses its airworthiness and is not available for operations
2. **Maintenance ground time minimization:** The ground time associated with the maintenance slots should be minimized. For example, it is more efficient to assign a task with a required duration of execution 15 h to a flexible maintenance slot with a duration of 16 h rather than to a slot with a duration of 20 h.
3. **Schedule stability:** A schedule change occurs when the aircraft registration assigned to a flexible maintenance slot in the existing schedule, r_{mor} , changes because of the arrival of a corrective task or an updated RUL prediction. It is preferred to avoid having schedule changes as we move towards the day of operations, i.e., a schedule change on the day of operations is more costly than a schedule change 10 days ahead.
4. **Task utilization:** The final objective is to plan tasks at the optimal moment in time. A widely used metric that airlines use in order to quantify the efficiency of task scheduling is task interval utilization. The task interval utilization, TU_g , $g \in G_r$, for the different types of tasks is calculated as follows:

$$TU_g = \begin{cases} \frac{Start_m - Arrival_g}{Due_g - Arrival_g}, & \text{if } g \in G_{corr}^r \cup G_{prev}^r \\ \frac{Start_m - Arrival_g}{EoL_{u_r} - Arrival_g}, & \text{if } g \in G_{progn}^r \end{cases} \quad (5)$$

$Start_m$ is the start date of maintenance slot m , $Arrival_g$ is the creation date of maintenance task g , Due_g is the due date of task g , and EoL_g is the true end-of-life of the monitored component u_r that corresponds to the prognostics-driven task g . For the preventive tasks and the prognostics-driven tasks, the objective is to schedule them as close as possible to the due date and the end-of-life of the component, respectively, in order to minimize maintenance interventions in the future, meaning that a high value of TU_g should be achieved. On the contrary, the corrective tasks should be resolved as soon as possible for quality reasons, so a low TU_g is expected.

3. Maintenance Scheduling of Prognostics-Driven Tasks

We formulate the decision-making process for the execution of prognostics-driven tasks as a partially observable Markov decision process (POMDP), which is solved using a modified version of the POMCP algorithm. Readers are referred to [13] for a detailed explanation of the POMDP formulation and the POMCP algorithm. The general concept is that the POMCP algorithm constructs a search binary-decision tree of state histories or beliefs for every considered component (associated prognostics-driven task). An example of such a tree, for a component with two hidden health states (Healthy, Degrading) and one evident state (Fail) is visualized in Figure 2. The algorithm uses the following inputs:

1. The RUL predictions for every component, which for the purposes of this study are assumed to follow the normal distribution, i.e., at time $t = T_n$, $\overline{L_{u_r}(T_n)} \sim N(\mu_{u_r}, \sigma_{u_r})$.
2. The available maintenance slots, $m \in M$.
3. The average daily aircraft utilization, δ_r .
4. The desired planning horizon.
5. The component maintenance cost at time $t = T_n$, $C_{u_r}(T_n)$, which is formulated as a combination of the corrective, $C_{u_r}^{corr}$, and the preventive maintenance cost, $C_{u_r}^{prev}$, of the component u_r as follows:

$$C_{u_r}(T_n) = C_{u_r}^{corr} \times P_{u_r}^{fail}(T_{n+1}) + C_{u_r}^{prev} \times [(1 - P_{u_r}^{fail}(T_{n+1})) + \frac{\mathbb{E}[(\overline{L_{u_r}(T_n)}) | TSI_{u_r}(T_n)] - TSI_{u_r}(T_n)}{MTBF_{u_r}}] \quad (6)$$

where $TSI_{u_r}(T_n) = n \cdot \delta_r$ corresponds to the elapsed time from the installation of the component until $t = T_n$, $MTBF_{u_r}$ is the Mean Time Between Failures for the specific component, and $P_{u_r}^{fail}(T_{n+1})$ is the probability that the component will fail until the next decision epoch $n + 1$. The probability of failure of the component is calculated as follows:

$$P_{u_r}^{fail}(T_{n+1}) = \begin{cases} 0 & \text{if Health State = Healthy} \\ P(\overline{L_{u_r}(T_n)} < \delta_r) = \Phi(\frac{\delta_r - \mu_{u_r}}{\sigma_{u_r}}) & \text{if Health State = Degrading} \end{cases} \quad (7)$$

where $\Phi(\cdot)$ corresponds to the CDF of the normal distribution.




Figure 2. Monte Carlo component tree.

We assume that the predictions from the prognostics model are available every day, so the decision epoch n corresponds to day n of the planning horizon. Accordingly, the tree is organized in n alternating layers of belief b_{T_n} and action nodes, a_{T_n} where each layer corresponds to a day of the planning horizon. Each node is characterized by the number of visits N , which counts the number of times this node has been visited, and a value V , which captures the average estimated return of all simulations when starting from this node.

Starting from the root node, which corresponds to the belief of the maintenance planner at the present day $t = T_0$, we perform multiple iterations to explore the different branches of the tree, or alternatively, the different scenarios regarding the health state of the component. As a result of the chosen action at each action node, the maintenance planner receives a total discounted accumulated return $R_{T_n}^u = \sum_{t=T_n}^{T_{end}} \gamma r_t^u$, where γ is a discount factor and r_t^u is the difference of maintenance cost between two consecutive decision epochs and can be formulated as follows:

$$r_t^u = C(t-1) - C(t) \quad (8)$$

This formulation of the reward function is intended to capture the additional maintenance cost savings or losses that can be incurred because of the decision of the maintenance planner to postpone the maintenance of the component for one additional day. Then, the value function, which can be used to assess the quality of action α at time $t = T_n$ can be written as:

$$V_\alpha^u(b_{T_n}) = \mathbb{E}_\pi[R_{T_n}^u | b_{T_n}] \quad (9)$$

and corresponds to the expected return that will be earned over the planning horizon $[T_n, T_{end}]$, starting from belief state b_{T_n} . When all the simulations are complete, the maintenance planner selects the action node with the greatest value function:

$$\hat{a} = \operatorname{argmax}_{a_t}(V^u(b_{t a_t})) \quad (10)$$

Finally, when a new prediction from the prognostics has been received, we prune the tree at the belief node determined by the received observation. This specific belief node becomes the new root node of the tree and, as such, all the other belief nodes are now impossible.

4. Schedule Optimization Models

4.1. MILP Scheduling Model

In this section, the model developed by [11] is briefly discussed. To implement the scheduling objectives described in Section 2.4, the objective function is formulated as a weighted sum of costs, where the costs correspond to the different planning objectives. The decision variables are described in Table 1, whereas the mathematical formulation of the objective function is as follows:

$$\begin{aligned} \text{Min} \quad & \sum_{g \in G_{\text{Due}}} T_{g,m=0} \cdot W_{\text{DUE}} \cdot C_{\text{Type},g} + \\ & \sum_{m \in M} \left(AC_{r_{\text{mor}},m} + \sum_{r \in R_m \setminus r=r_{\text{mor}}} AC_{r,m} \cdot (1 + W_{\text{RES},m}) \right) \cdot \text{Duration}_m \cdot W_{\text{GROUND}} + \\ & \sum_{m \in M} \left(\sum_{g \in G_{\text{prev}}} T_{g,m} \cdot W_{\text{INTERVAL prev},g,m} + \right. \\ & \left. \sum_{g \in G_{\text{progn}}} T_{g,m} \cdot W_{\text{INTERVAL progn},g,m} + \sum_{g \in G_{\text{corr}}} T_{g,m} \cdot W_{\text{INTERVAL corr},g,m} \right) \cdot C_{\text{Type},g} \end{aligned} \quad (11)$$

The defined weights are captured in Table 2 as follows.

Table 1. Defined decision variables for MILP model.

Decision Variable	Description
$T_{g,m}$	Binary, 1 if task g is assigned to slot m
$AC_{r,m}$	Binary, 1 if aircraft r is assigned to slot m

Table 2. Defined weights for the MILP model.

Objectives	Weights	Values
Timely task execution	W_{DUE}	10^8
Maintenance ground time minimization	W_{GROUND}	10^5
Task utilization	W_{INTERVAL}	[0–1]

The weight related to the objective of schedule stability is formulated as a function of the number of days of notice, as visualized in Figure 3, and formulated as follows:

$$W_{\text{RES},m} = \max\left(\eta \cdot \delta - \frac{\eta}{\delta} \cdot (\text{Start}_m - \text{Current Date}), 0\right) \quad (12)$$

where δ is the defined days of notice, or in other words, the desired period in which schedule changes should be prevented.




Figure 3. Rescheduling cost as a function of days of notice [11].

$C_{Type,g}$ corresponds to the task criticality based on the type of the task and its impact on aircraft airworthiness. The values for $C_{Type,g}$ are defined in Table 3.

Table 3. Weights for task type.

Task Type	$C_{Type,g}$
Preventive	4
Corrective	[1–4]
Prognostics-driven	1

To create a feasible schedule that adheres to the regulatory and operational constraints of a real airline environment, the set of constraints (1)–(4) are considered. A more elaborate formulation of the model, constraints, and parameters is provided in [11]. The MILP model was solved using GUROBI commercial solver.

4.2. Deep Reinforcement Learning Scheduling Model

In this section, a brief overview of the DRL algorithm developed in [13] is presented. Readers are referred to the corresponding work for a detailed explanation of the state features, reward function, and scheduling algorithm. The general concept is that the reinforcement learning algorithm considers each maintenance opportunity m at a time and decides which aircraft to schedule, subject to constraints (1)–(4). More specifically, based on the state of the environment s_m , the algorithm selects the action $a_r \in \mathcal{A}_r^m$ with the highest value function Q . Each action a_r corresponds to the decision to perform maintenance on aircraft r .

A deep Q-network (DQN) is developed, consisting of five fully connected layers, with one input layer, three hidden layers, and one output layer. The structure of the DQN is illustrated in Figure 4.




Figure 4. DQN structure.

In order to prevent having a varying input and output layer size due to the different number of aircraft with open tasks at every decision point, we choose to always consider at every decision point the 10 most critical aircraft in terms of remaining maintenance opportunities, i.e., we sort all the aircraft in ascending order according to the remaining opportunities until their first task goes due and we select the first 10 from this sorted list. This results in a number of 134 nodes for the input layer and 12 nodes for the output layer. Additionally, each hidden layer consists of 64 neurons. We use the ReLu activation function for the input and the hidden layers and the linear activation function for the output layer.

5. Computational Experiments

5.1. Case Study

In this section, the two scheduling models will be applied to three different maintenance scenarios for different aircraft fleet sizes with data provided by a major European airline. The data are spread over a period of 5 months of airline operations. According to the requirements defined in MPD and MEL, for each aircraft there is a list of open preventive and corrective maintenance tasks which are updated on a daily basis. The execution of these tasks can be performed in specified maintenance opportunities. Each maintenance opportunity has specific available resources, which vary over time. Moreover, the available workforce in each maintenance opportunity is organized in different skills. Data about the (arrival of) corrective and preventive maintenance tasks, the maintenance slots schedule, and the available resources are provided by the airline.

Furthermore, for 10 aircraft from the considered fleets, we simulate additionally a total of 250 prognostics-driven tasks. More specifically, we assume that each aircraft has 25 monitored components with predictable RULs which are updated on a daily basis and are assumed to follow the normal distribution. The working state of the components included in the 10 aircraft is different, i.e., every component has a different true RUL.

The predictions are simulated by applying the Support Vector Regression (SVR) prognostics model developed in [15] on the C-MAPSS dataset [16] for turbo-fan engines, assuming that the time cycles used in the C-MAPSS dataset correspond to flight cycles (FCs). Following the approach described in [17], the obtained predictions are then organized in four clusters according to the prediction accuracy and uncertainty, described by MAE and standard deviation σ in FCs, respectively:

- **Cluster #1:** $MAE \sim 8.32$ and $\sigma \sim 12.24$;
- **Cluster #2:** $MAE \sim 13.89$ and $\sigma \sim 19.69$;
- **Cluster #3:** $MAE \sim 23.52$ and $\sigma \sim 27.35$;
- **Cluster #4:** $MAE \sim 37.28$ and $\sigma \sim 40.93$.

We then assume that for the 28% of the components (7 components per aircraft in total) we obtain continuously updated RUL predictions belonging to cluster 1, for 24% to cluster 2 (6 components per aircraft in total), for 24% to cluster 3 (6 components per aircraft in total), and for 24% to cluster 4 (6 components per aircraft in total).

Moreover, the corrective and preventive replacement cost of every monitored component is set to $C_u^{Corr} = 25,000$ and $C_u^{Prev} = 10,000$, respectively. The magnitude of the cost values was driven by [18], where average historical values of true preventive and corrective repair costs were used.

In order to simulate the dynamic process of the arrival of new corrective and preventive tasks, and also the update of RUL predictions for the monitored components, we implement a rolling horizon approach (see Figure 5). A maintenance schedule is generated for a fixed time window. Afterward, the planning horizon shifts one day ahead, where new tasks may arrive and/or new RUL predictions are obtained. The scheduling algorithms produce then a new and feasible schedule for the intended time window, while we choose to minimize the number of schedule changes in the next 3 days (highlighted grey area in Figure 5). The choice of 3 days is driven by the airline practice, as aircraft allocation to flight usually occurs 3 days before the day of operations. The same process repeats until the end of the

planning horizon is reached. It is noted that no scheduling opportunities beyond the end of the planning horizon are considered.




Figure 5. Rolling horizon approach [11].

5.2. Assumptions

In our case study, we adopt the following assumptions:

1. Aircraft utilization is known and constant. The daily aircraft utilization δ_r , $r \in R$ is set to 15 FHs or 4 FCs, according to historical aircraft utilization values of an airline.
2. RUL predictions for every monitored component are assumed to follow the normal distribution.
3. The prognostics-driven tasks correspond to components that are not critical for the safe operation of the aircraft.

5.3. Results Analysis

The scheduling models are going to be compared and evaluated on the basis of the four maintenance planning objectives described in Section 2.4. Table 4 summarizes the obtained results after applying the MILP and the DRL scheduling models on different maintenance scenarios for the different aircraft fleet sizes.

Table 4. Scheduling results for the MILP and DRL algorithm.

Maintenance scenario		#1		#2		#3	
Fleet size (# aircraft)		15		26		40	
# Tasks		899		1007		1624	
Timely task execution	Model	MILP	DRL	MILP	DRL	MILP	DRL
	Due tasks (Prognostics-driven)	4	5	6	6	3	4
	Due tasks (Corrective/Preventive)	0	0	0	0	2	0
Maintenance ground time minimization	Used Slots	163	169	184	187	289	292
	Total ground time (h)	1972.24	2100.13	2235.75	2338.56	5332.68	5444.92
	Average slot duration	12.09	12.42	12.15	12.50	19.11	19.30
Schedule stability	Schedule changes	6	12	10	17	11	20
Task utilization (%)	Preventive tasks utilization	75.8	75.6	75.9	74.1	78.3	77.2
	Corrective tasks utilization	51.4	42.9	55.4	53.6	52.3	49.1
	Prognostics-driven tasks utilization	38.2	72.7	42.2	71.1	41.1	74.1

With respect to timely execution, both models achieve approximately the same results, with the MILP algorithm performing better in scenario #1 and the DRL algorithm achieving better results in scenario #3. In all scenarios, the DRL algorithm requires more maintenance slots, which is translated as a higher allocated ground time for maintenance. However,

as an exchange for the increased ground time, the DRL algorithm performs significantly better than the MILP algorithm when it comes to the RUL exploitation of the monitored components, utilizing on average across all scenarios 72.6% of their RUL, compared to 40.5% achieved by the MILP algorithm. Moreover, even though the DRL algorithm uses more maintenance slots, the average duration of the chosen slots is almost similar to the average duration of slots used by the MILP algorithm.

The MILP algorithm, being more conservative with respect to the scheduling of the prognostics-driven tasks, induces fewer last-minute changes to the maintenance schedule than the DRL algorithm in all scenarios. Furthermore, the DRL algorithm performs slightly better than the MILP algorithm when scheduling the corrective tasks, whereas the MILP algorithm performs slightly better when scheduling the preventive tasks.

Finally, the computational time requirements of both models are summarized in Table 5. These computational times were obtained using an Intel Core i5 processor 9300, with 16GB RAM and NVIDIA 1660Ti GPU. In all scenarios, the computation time needed by the DRL algorithm to produce the maintenance schedule for every day of the planning horizon is lower compared to the MILP algorithm, with the DRL being as high as approximately 65% faster. The former highlight the scalability of the DRL algorithm when more aircraft and tasks are considered. Nevertheless, both models are proven to be suitable for quasi real-time decision making, as required in an airline maintenance environment.

Table 5. Computational time of the MILP and DRL algorithm.

Maintenance scenario	#1		#2		#3	
Fleet size	15		26		40	
Tasks	899		1007		1624	
Model	MILP	DRL	MILP	DRL	MILP	DRL
Computational time (s)	13.84	5.9	15.66	6.3	25.32	7.1

From the results, it can be concluded that both presented models are suitable for scheduling in an airline environment, reflecting different maintenance scheduling strategies. The MILP model assigns more weight to maintenance decisions targeting to minimization of ground time and increased schedule stability, whereas the DRL model assigns more weight to the exploitation of RUL of the monitored components. This observed trade-off provides also valuable insights with respect to aircraft scheduling in a CBM environment. As shown from the results, in order to leverage the benefits of the CBM strategy with respect to a high RUL exploitation of the components, potentially more aircraft visits to the hangar for maintenance will be required. This is a direct consequence of the uncertainty included in the RUL predictions from the prognostics models. This means that for an effective transition towards a CBM approach, the airlines should consider adding more flexibility to their maintenance schedule. However, in the long term, this would ultimately lead to less replacement/repairs of the monitored components and lower inventory-related costs.

6. Conclusions

In this paper, a comparison between a MILP and a DRL model for the maintenance scheduling of an aircraft fleet in a CBM context was presented. The RUL prognostics, are updated on a daily basis with new sensor measurements, and are characterized by uncertainty that follows the normal distribution. On top of that, also the list of preventive and corrective maintenance tasks is continuously updated. Both scheduling models take into account the list of different types of maintenance tasks, along with available maintenance slots, the available resources and the existing maintenance schedule, to produce the maintenance schedule of the aircraft fleet using a rolling horizon approach. The overarching goal is to prevent tasks going due, while at the same time, ensuring high fleet availability, schedule stability and efficient task interval utilization.

The performance of the investigated models was evaluated on three real maintenance scenarios for different aircraft fleet sizes with data provided by our partner airline, enriched with simulated data for the prognostics-driven tasks from the C-MAPSS dataset. The results show that both models have similar performance with respect to timely task execution. The DRL algorithm manages to schedule more efficiently the prognostics-driven tasks, achieving a higher RUL exploitation for the monitored components. On the other hand, the MILP algorithm induces less maintenance ground time and requires last-minute changes in the maintenance schedule. As such, the choice of which model to use relates to the objective that the airline considers of higher importance. Finally, considering quasi real-time requirements for applications of scheduling models in a real airline environment, both models achieve computational times below one minute. However, the results highlight the capability of DRL to have an increased computational efficiency that stays unaffected of the problem size and considered variables.

Future work may focus on examining the robustness of both models in different types of RUL prognostics distributions, with different types of prediction accuracy. Furthermore, it would be interesting to investigate the implementation of a hybrid approach, i.e. using a MILP model to facilitate the offline training of the DRL model, as this would help to improve the quality of the solution.

Author Contributions: Conceptualization, I.T., P.J.v.K. and B.F.S.; methodology, I.T. and P.J.v.K.; software, I.T. and P.J.v.K.; formal analysis, I.T. and P.J.v.K.; investigation, I.T. and P.J.v.K.; resources, B.F.S.; data curation, I.T. and P.J.v.K.; writing—original draft preparation, I.T.; writing—review and editing, I.T., P.J.v.K. and B.F.S.; visualization, I.T. and P.J.v.K.; supervision, B.F.S.; project administration, B.F.S.; funding acquisition, B.F.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the European Union’s Horizon 2020 research and innovation program under the REMAP project, grant number 769288.

Data Availability Statement: The data used in this study are not available due to 3rd party restrictions.

Conflicts of Interest: The authors declare no conflict of interest.

Nomenclature

The following notations are used in this manuscript:

g	Open maintenance tasks that require execution
r	Aircraft indicator
m	Maintenance slot indicator
w	Workforce skill indicator
u	Monitored components
$g \in G$	Set of task groups $G = G_{\text{prev}} \cup G_{\text{corr}} \cup G_{\text{progn}}$
$g \in G_{\text{prev}}$	Subset of preventive tasks ($G_{\text{prev}} \subset G$)
$g \in G_{\text{corr}}$	Subset of corrective tasks ($G_{\text{corr}} \subset G$)
$g \in G_{\text{progn}}$	Subset of prognostics-driven tasks ($G_{\text{progn}} \subset G$)
$g \in G^r$	Subset of tasks for aircraft r ($G^r \subset G$)
$r \in R_m$	Set of aircraft of the aircraft type of slot m
$m \in M$	Set of maintenance slots within the current schedule window
$m \in M_{\text{Fixed}}$	Subset of slots which the aircraft is fixed ($M_{\text{Fixed}} \subset M$)
$m \in M_{\text{Flexible}}$	Subset of slots which the aircraft is allowed to change ($M_{\text{Flexible}} \subset M$)
$w \in W$	Set of workforce skills
$C_{u_r}^{\text{prev}}$	Preventive maintenance cost for the monitored component u_r
$C_{u_r}^{\text{corr}}$	Corrective maintenance cost for the monitored component u_r

GR_g^w	Required workhours of skill w for the execution of task g
GR_m^w	Available workhours of skill w on maintenance slot m
$Duration_g$	Required duration of task g
$Duration_m$	Duration of slot m
$Start_m$	Start date of maintenance slot m
$Arrival_g$	Creation date of maintenance task g
Due_g	Due date of task g
EoL_g	True end-of-life of component u_r that corresponds to prognostics-driven task g
$MTBF_{u_r}$	Mean Time Between Failures for monitored component u_r
δ_r	daily utilization rate of aircraft r

References

1. Sprong, J.; Jiang, X.; Polinder, H. Deployment of Prognostics to Optimize Aircraft Maintenance—A Literature Review. *J. Int. Bus. Res. Mark.* **2020**, *5*, 26–37. [\[CrossRef\]](#)
2. Deng, Q.; Santos, B.; Curran, R. A practical dynamic programming based methodology for aircraft maintenance check scheduling optimization. *Eur. J. Oper. Res.* **2020**, *281*, 256–273. [\[CrossRef\]](#)
3. Lagos, C.; Delgado, F.; Klapp, M. Dynamic Optimization for Airline Maintenance Operations. *Transp. Sci.* **2020**, *54*, 998–1015. [\[CrossRef\]](#)
4. Chansombat, S.; Pongcharoen, P.; Hicks, C. A mixed-integer linear programming model for integrated production and preventive maintenance scheduling in the capital goods industry. *Int. J. Prod. Res.* **2019**, *57*, 61–82. [\[CrossRef\]](#)
5. Pham, D.N.; Klinkert, A. Surgical case scheduling as a generalized job shop scheduling problem. *Eur. J. Oper. Res.* **2008**, *185*, 1011–1025. [\[CrossRef\]](#)
6. Qin, Y.; Wang, Z.X.; Chan, F.T.; Chung, S.H.; Qu, T. A mathematical model and algorithms for the aircraft hangar maintenance scheduling problem. *Appl. Math. Model.* **2019**, *67*, 491–509. [\[CrossRef\]](#)
7. Erdem, E.; Qu, X.; Shi, J. Rescheduling of elective patients upon the arrival of emergency patients. *Decis. Support Syst.* **2012**, *54*, 551–563. [\[CrossRef\]](#)
8. Vali-Siar, M.; Gholami, S.; Ramezani, R. Multi-period and multi-resource operating room scheduling under uncertainty: A case study. *Comput. Ind. Eng.* **2018**, *126*, 549–568. [\[CrossRef\]](#)
9. Ballestin, F.; Perez, A.; Quintanilla, S. Scheduling and rescheduling elective patients in operating rooms to minimise the percentage of tardy patients. *J. Sched.* **2019**, *22*, 107–118. [\[CrossRef\]](#)
10. Liu, X.; Sun, Q.; Ye, S.; Yildirim, M. Optimal multi-type inspection policy for systems with imperfect online monitoring. *Reliab. Eng. Syst. Saf.* **2021**, *207*, 107335. [\[CrossRef\]](#)
11. van Kessel, P.J.; Freeman, F.C.; Santos, B.F. Airline maintenance task rescheduling in a disruptive environment. *Eur. J. Oper. Res.* **2022**, *in press*. [\[CrossRef\]](#)
12. Andrade, P.; van Silva, C.; Ribeiro, B.; Santos, B. Aircraft Maintenance Check Scheduling Using Reinforcement Learning. *Aerospace* **2021**, *8*, 113. [\[CrossRef\]](#)
13. Tseremoglou, I.; Santos, B. Condition-Based Maintenance Scheduling of an Aircraft Fleet Under Partial Observability: A Deep Reinforcement Learning Approach. *WorkingPaper* 2022, *unpublished*.
14. Silver, D.; Veness, J. Monte-Carlo planning in large POMDPs. In *Advances in Neural Information Processing Systems*; Vancouver, BC, Canada, 2010; pp. 1–9.
15. Bieber, M.; Verhagen, W.; Santos, B. An Adaptive Framework For Remaining Useful Life Predictions of Aircraft Systems. *PHM Soc. Eur. Conf.* **2021**, *6*, 60–70. [\[CrossRef\]](#)
16. Saxena, A.; Goebel, K. Turbofan Engine degradation simulation dataset. In *NASA Ames Prognost Data Reposit*; 2008; pp. 878–887.
17. Tseremoglou, I.; Bieber, M.; Santos, B.; Floris, F.; van Kessel, P. The Impact of Prognostic Uncertainty on Condition-Based Maintenance Scheduling: An Integrated Approach. In *Proceedings of the AIAA Aviation Forum*, Chicago, IL, USA, 27 June–1 July 2022; American Institute of Aeronautics and Astronautics Inc.: Reston, VA, USA, 2022. [\[CrossRef\]](#)
18. Freeman, F.; van Kessel, P.; Verhagen, W. Age and Condition-Based Preventive Replacement Timing for Periodic Aircraft Maintenance Checks. In *Proceedings of the 6th European Conference of the Prognostics and Health Management Society 2021*, Virtual Event, 28 June–2 July 2021; pp. 151–161.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.