Received 11 March 2025
Revised 22 December 2025
8 March 2026
10 March 2026
Accepted 11 March 2026

Leveraging simulation and robust decision-making to optimize military asset maintenance and workforce management

Hasan Hüseyin Turan, Sanath Darshana Kahagalage and Sondoss El Sawah
School of Systems and Computing, University of New South Wales, Canberra, Australia

Abstract

Purpose – This study develops an integrated simulation and robust decision-making framework to support long-term military asset maintenance and workforce planning under deep uncertainty. The objective is to evaluate how alternative maintenance workforce strategies perform across a wide range of plausible future conditions where key system parameters cannot be reliably estimated.

Design/methodology/approach – A discrete-event simulation model is developed to represent maintenance operations, workforce competency progression, certification dynamics and resource constraints within a military maintenance system. The model is integrated with an exploratory modelling and analysis (EMA) framework implementing robust decision-making (RDM). Predefined workforce recruitment and periodic maintenance strategies are evaluated across ensembles of scenarios generated using Latin Hypercube sampling, followed by robustness assessment and scenario discovery analysis.

Findings – Results indicate that strategies combining lower workforce recruitment with higher maintenance intensity can outperform intuitively preferred high-resource strategies when evaluated under deep uncertainty. The analysis identifies workforce separation rates – particularly among coordinator and verifier roles – as key drivers of system vulnerability. The framework also reveals critical uncertainty thresholds that may act as early warning indicators of degradation in maintenance performance.

Research limitations/implications – The study focuses on a single asset class and workforce type to enable detailed behavioural representation. Future research may extend the framework to multiple asset portfolios, adaptive strategy design and dynamic policy adjustment.

Practical implications – The proposed approach provides defence planners with a transparent decision-support tool for evaluating maintenance and workforce policies when historical data are limited, enabling identification of robust strategies and workforce indicators that should be monitored to sustain long-term readiness.

Originality/value – This study presents an integrated framework combining simulation-based asset management modelling with robust decision-making to evaluate maintenance and workforce strategies for military systems under deep uncertainty. By assessing strategies across many scenarios using robustness metrics, the framework supports more resilient long-term sustainment planning.

Keywords Military maintenance workforce, Robust decision-making, Exploratory modelling and analysis, Deep uncertainty, Periodic maintenance, Simulation

Paper type Research article

1. Introduction

The military maintenance workforce operates within a uniquely constrained, closed-loop environment in which promotions, internal training pathways, and asset-dependent certification requirements shape workforce availability and progression (Withers *et al.*, 2021). These interdependencies distinguish defence maintenance systems from typical

Journal of Defense Analytics and Logistics
Vol. 10 No. 1, 2026
pp. 20-40
Emerald Publishing Limited
e-ISSN: 2399-6447
p-ISSN: 2399-6439
DOI 10.1108/JDAL-03-2025-0007

industrial settings and motivate the need for modelling tools capable of representing their structural complexity.

Maintenance is fundamental to sustaining asset readiness, reliability, and availability, particularly for high-value military capabilities where system failure carries substantial operational risk. Organisations use either reactive maintenance—conducted after failures—or periodic maintenance, which proactively schedules inspections and service intervals to reduce unplanned downtime and improve predictability. The effectiveness of periodic maintenance, however, depends critically on the availability, competency progression, and utilisation of the skilled maintenance workforce.

Despite extensive literature on maintenance scheduling (Jafar-Zanjani *et al*., 2022; Shahmoradi-Moghadam *et al*., 2021; Zhang *et al*., 2023) and on workforce planning (De Bruecker *et al*., 2015), these streams have largely evolved in isolation. Existing studies optimise maintenance activities or workforce structures separately and therefore overlook their joint effects in contexts where skilled personnel are both essential resources and binding constraints (Safaei *et al*., 2011). To our knowledge, there is no integrated modelling framework that simultaneously represents (1) periodic maintenance scheduling, (2) workforce planning and progression, and (3) the deep uncertainties affecting both—particularly within defence environments. This gap is consequential, as modern acquisition programs introduce new asset classes for which historical data on technician availability, separation, and training requirements may be sparse or nonexistent.

Deep uncertainty further complicates long-term planning, as stakeholders often lack agreement on appropriate model structures, parameter values, or outcome preferences (Walker *et al*., 2010). While simulation supports detailed what-if analysis, single-scenario assessments provide limited insight when key parameters—such as workforce separation rates or maintenance resource reliability—are themselves uncertain. Robust Decision-Making (RDM) and exploratory modelling provide a systematic means to evaluate strategies across a wide range of plausible futures rather than relying on best-estimate assumptions.

This paper examines the interaction between periodic maintenance scheduling and workforce planning under deep uncertainty using an integrated simulation-based approach. We evaluate predefined maintenance–workforce strategies across a large ensemble of futures and identify the strategy that remains most robust to uncertainty. We further analyse the conditions under which even the most robust strategy may underperform, providing insight into critical uncertainties and early-warning indicators.

The contributions of this study are twofold. (1) We develop a detailed simulation model—the Asset Management and Maintenance Model (AMMM) —that represents key defence workforce dynamics, including competency progression, certification decay, mentor-dependent task execution, and failures of maintenance resources. Such features, to the best of our knowledge, have not been jointly incorporated in prior defence maintenance simulation studies. (2) We integrate AMMM with an exploratory modelling and analysis framework to evaluate maintenance and workforce strategies when historical data are limited or uncertain. This enables the identification of a robust strategy as well as the discovery of critical uncertain parameters whose ranges define vulnerable futures.

In addition, the study provides managerial insight into how predefined maintenance schedules and workforce recruitment policies interact under skilled-labour constraints, and how these interactions influence readiness, workload balance, and maintenance timeliness. The global sensitivity and vulnerability analyses offer practical tools for monitoring critical uncertainties and anticipating conditions under which strategic adjustments may be required.

The remainder of this paper is organised as follows. Section 2 reviews relevant literature. Section 3 presents the AMMM architecture and formulates the decision problem. Section 4 describes the exploratory modelling and robust decision-making approach. Section 5 reports the results, including strategy comparison and vulnerability analysis. Section 6 concludes with implications for defence maintenance planning.

2. Literature review

Workforce planning has been extensively studied across engineering, logistics, and defence domains, with emphasis on recruitment pipelines, ageing workforce challenges, skill development, retention, and organisational mobility (Kafiabad *et al.*, 2020). These studies highlight the strong interdependencies that shape workforce availability—such as training lead times, learning curves, and competency progression—which are particularly critical in technical maintenance environments. In military contexts, delays in developing or losing skilled personnel translate directly into reduced maintenance capacity and diminished asset readiness.

Parallel to this, a substantial body of work examines maintenance modelling and optimisation. Research on periodic and preventive maintenance planning (Jafar-Zanjani *et al.*, 2022; Mansour, 2011) addresses reliability improvement, cost minimisation, and scheduling efficiency across operational settings including flowshops (Perez-Gonzalez *et al.*, 2020), single- and parallel-machine systems (Yazdani *et al.*, 2023; Sbihi and Varnier, 2008; Liao and Chen, 2003; Shen and Zhu, 2019), and multi-component repairable systems (Dong *et al.*, 2019; Fu and Wang, 2022). Although these models provide valuable insights into optimal maintenance timing and resource allocation, they typically assume fixed or aggregate resource availability. As a result, they do not capture the underlying workforce dynamics—turnover, certification loss, training throughput—that constrain real-world maintenance execution.

Despite rich literature on workforce planning and maintenance modelling individually, their integration remains limited. Workforce planning models rarely simulate the operational consequences of maintenance demand, while maintenance models seldom incorporate the detailed workforce processes that determine actual throughput. This disconnect is especially prominent in defence settings, where mission-critical assets depend on specialist technicians whose availability is shaped by complex career structures and certification requirements (Kafiabad *et al.*, 2020).

2.1 A limited literature in military maintenance workforce planning and maintenance scheduling

Scant literature addresses the dual challenge of maintenance scheduling and workforce planning in military contexts. Safaei *et al*. (2011) examined preventive maintenance for military aircraft under skilled-worker constraints, showing that workforce scarcity is a dominant bottleneck. Their time-indexed mixed-integer programming model sought to maximise fleet availability subject to restricted technician capacity. Zhang *et al*. (2023) developed a heuristic scheduling approach to maximise fleet readiness in a similar setting.

Workforce planning in military environments has also been explored through joint optimisation formulations (Turan *et al.*, 2021a,b) and simulation-based planning tools (Horn *et al.*, 2016). Akl *et al*. (2022) considered combined strategic workforce planning and preventive maintenance scheduling, but again with a narrow workforce representation. Overall, publicly available literature either treats maintenance scheduling or workforce planning in isolation, or represents workforce capacity in overly aggregated terms. Little work examines how maintenance workforce capability and planning jointly influence the ability to execute critical mission maintenance. These studies also do not address deep uncertainty, which is prominent in defence environments and commonly treated in non-military settings (Weaver *et al.*, 2013; Lempert *et al.*, 2006; Kwakkel *et al.*, 2016).

Given that skilled technicians are both the primary resource and primary constraint in military workshops, existing optimisation-focused studies offer limited insight into how periodic maintenance and workforce planning interact when uncertainties in separations, training throughput, and maintenance resource reliability are significant.

2.2 Conventional decision support methods are not appropriate to handle the condition of deep uncertainty

Traditional deterministic or probabilistic decision-support methods are ill-suited for environments characterised by deep uncertainty, where model structure, parameter values, and stakeholder preferences cannot be agreed upon or reliably estimated (Stanton and Roelich, 2021; Morgan *et al.*, 1990). Such approaches risk overconfidence in point predictions and may lead to strategies that perform poorly when future conditions deviate from expectations.

Decision-making under uncertainty is often framed around two paradigms (Weaver *et al.*, 2013). The *Predict-then-Act* paradigm seeks optimal solutions based on single best-estimate assumptions or probabilistic distributions (Morgan *et al.*, 1990). In contrast, the *Seek Robust Solutions* paradigm identifies strategies that perform adequately across a wide range of plausible futures, revealing vulnerabilities and guiding the design of regret-minimising policies (Lempert *et al.*, 2006). Robust decision-making (RDM), which adopts the latter paradigm, is especially valuable in defence contexts with long planning horizons and limited historical data.

2.3 The model-based decision-making frameworks to confront deep uncertainty

Model-based decision frameworks such as RDM use exploratory modelling to examine many plausible futures, assess strategy performance, and identify critical uncertainties that drive system vulnerabilities (Weaver *et al.*, 2013; Hall, 2007). RDM has been widely applied in defence (Lempert, 2003), climate and water systems (Weaver *et al.*, 2013; Kwakkel *et al.*, 2016), and healthcare (Scholte *et al.*, 2023). It emphasises iterative vulnerability analysis, robustness metrics, and participatory structuring of the decision problem, making it well suited to planning under deep uncertainty.

Motivated by these principles, this study evaluates alternative combinations of periodic maintenance scheduling and workforce planning using RDM. We identify the most robust strategy across a wide uncertainty space, perform vulnerability analysis to determine the parameters and ranges associated with underperformance, and highlight the critical factors that decision-makers must monitor to sustain effective long-term maintenance capability.

3. Simulation architecture and decision model formulation

This section describes the major components of the maintenance activity simulation (Section 3.1) and the detailed formulation of the maintenance and workforce planning decision problem (Section 3.2).

3.1 The asset management and maintenance model (AMMM)

The architecture of AMMM (Figure 1) integrates maintenance system assets and their schedules, maintenance activities and their schedules, and workforce planning decisions and schedules. These are simulated and evaluated in the integrated AMMM environment.

The core components of the model consist of several interacting modules that collectively define system behaviour. The model includes entities such as assets and members of the maintenance workforce, including contractors, public service personnel, and military staff. Assets transition through a range of states during their lifecycle, including acquired, ready-for-issue, in use, preparing for maintenance, repairable or non-repairable, waiting for maintenance, undergoing maintenance, completed maintenance, awaiting storage, being transported to a facility, or expired. The movement and behaviour of these entities are governed by activities that represent the workflows and sequences of tasks each entity follows, where tasks form the smallest units of work and carry attributes such as duration and cost. Supporting these processes are resources, including storage facilities, maintenance equipment, and spare parts, which enable or constrain activity execution. The interactions among all components are further shaped by business rules, consisting of hard-coded logic that establishes the core operational dynamics through conditional statements, as well as flexible rules that users can modify as part of the experimentation setup.

The model offers a generic framework with a flexible structure. For instance, users have the ability to specify different maintenance schedules and maintenance workforce plans.

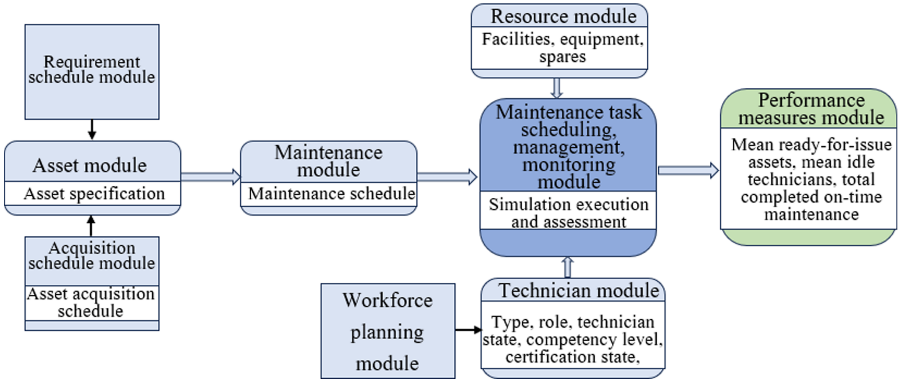

Figure 1. Architecture of the AMMM

The technician module represents a range of attributes that characterize technician behaviour and progression within the system. Each technician is assigned one of five competency levels—maintainer, verifier, inspector, coordinator, or manager—which apply across all workforce types. Technicians may enter the system at any level except manager, and they can progress to higher competency levels as they gain experience. Depending on maintenance needs, technicians perform different roles during tasks: they may serve as the primary technician, either fully authorised or in training; as a mentor, where an authorized technician guides a technician in training; or as a quality assurance technician who assesses and verifies task quality. Technicians also hold certification states linked to specific asset classes; upon entering the system, they are certified to work on one class, but they become uncertified if they do not complete the required number of tasks within a given time-frame. In addition, technicians are classified by employment type, including contractors, public service staff, and military personnel. At any given time, technicians may occupy one of several states—available, idle, busy, retired, or separated—reflecting their engagement and life-cycle status within the workforce.

The simulation considers only one type of maintenance workforce, represented by public service personnel, and focuses on a single asset type, referred to as the test asset. For each scenario, the yearly workforce separation fraction remains constant across all competency levels throughout the planning horizon. Similarly, for each strategy, the number of assets undergoing maintenance in each periodic event is fixed for the entire time horizon. The yearly failure rate and repair time of durable resources also remain constant within each scenario. Assets may enter the system only in the ready-for-issue (RFI), expired, or requiring-maintenance state; they cannot enter while already in the middle of maintenance. Assets are assumed to be acquired on their scheduled day without delay, and following maintenance, each asset transitions to the RFI state immediately and without delay.

3.2 Performance measures

Performance measures or outcome indicators encompass metrics that facilitate the evaluation of strategies under conditions of deep uncertainty based on the robustness criteria inferred from these measures. Regardless of the strategy employed, the objective remains to enhance available ready-for-issue (RFI) assets, minimize workforce idling, and boost on-time maintenance completion in accordance with the asset acquisition policy. Accordingly, the number of RFI assets, workforce idle time, and on-time maintenance completion are used as the three key performance measures for comparing alternative strategies.

To begin, we establish and introduce notations and terminology for various elements such as sets, decision variables, parameters, and collected values from the simulation runs. Note that a strategy is constructed using decision variables.

Sets and indices.

A: The set of assets with index $i \in A$

J: The set of technician's competency levels with index $j \in J$

K: The set of types of resources require for maintenance of assets with index $k \in K$

R: The set of durable maintenance resources indexed by r, where $r \in R$.

$\boldsymbol{W}^h$, $\boldsymbol{W}^m$, $\boldsymbol{W}^l$: Workforce plans representing high, medium, and low workforce requirements, respectively, for each technician competency level j (excluding managerial roles).

$\boldsymbol{M}^h$, $\boldsymbol{M}^m$, $\boldsymbol{M}^l$: Maintenance schedules representing high, medium, and low levels of maintenance intensity, respectively, defined by the number of periodic maintenance events per year and the number of assets undergoing maintenance. $t = 0, 1, \ldots, T-1, T$: Discrete time stage

Ω: A set of strategies

Decision variables.

$num\_e_i \in M*$ for $* \in \{h, m, l\}$: the number of periodic maintenance events per year for asset $i \in A$.

$num\_a_i \in M*$ for $* \in \{h, m, l\}$: the number of assets $i \in A$ undergoing maintenance during a periodic maintenance event. $\alpha^M \in W*$ for $* \in \{h, m, l\}$: Yearly recruitment rate of maintainers.

$\alpha^V \in W*$ for $* \in \{h, m, l\}$: Yearly recruitment rate of verifiers.

$\alpha^I \in W*$ for $* \in \{h, m, l\}$: Yearly recruitment rate of inspectors.

$\alpha^C \in W*$ for $* \in \{h, m, l\}$: Yearly recruitment rate of coordinators.

Parameters and inputs.

T: The end of the planning period

x: A strategy for $x \in \Omega$

Uncertain model parameters and their bounds.

s^M: The yearly separation fraction of Maintainers, where $s^M \in [\underline{s}^M, \overline{s}^M] \subset [0, 1]$, and $\underline{s}^M$ and $\overline{s}^M$ denote the lower and upper bounds of the separation fraction, respectively.

s^V: The yearly separation fraction of Verifiers, where $s^V \in [\underline{s}^V, \overline{s}^V] \subset [0, 1]$, and $\underline{s}^M$ and $\overline{s}^M$ denote the lower and upper bounds of the separation fraction, respectively.

s^I: The yearly separation fraction of Inspectors, where $s^I \in [\underline{s}^I, \overline{s}^I] \subset [0, 1]$, and $\underline{s}^I$ and $\overline{s}^I$ denote the lower and upper bounds of the separation fraction, respectively.

s^C: The yearly separation fraction of Coordinators, where $s^C \in [\underline{s}^C, \overline{s}^C] \subset [0, 1]$, and $\underline{s}^C$ and $\overline{s}^C$ denote the lower and upper bounds of the separation fraction, respectively.

s^{Ma}: The yearly separation fraction of Managers, where $s^{Ma} \in [\underline{s}^{Ma}, \overline{s}^C] \subset [0, 1]$, and $\underline{s}^{Ma}$ and $\overline{s}^{Ma}$ denote the lower and upper bounds of the separation fraction, respectively.

fr_r: Yearly failure rate of durable maintenance resource $r \in R$, where $fr_r \in \left[\underline{fr}_r, \overline{fr}_r\right] \subset [0, 1]$, and $\underline{fr}_r$ and $\overline{fr}_r$ denote the lower and upper bounds of the failure rate, respectively.

rt_r: Repair time of durable maintenance resource $r \in R$, where $rt_r \in [\underline{rt}_r, \overline{rt}_r]$, and $\underline{rt}_r$ and $\overline{rt}_r$ denote the lower and upper bounds of the repair time, respectively.

The analysis evaluates each predefined strategy through a set of performance measures that act as multi-criteria objectives. Let MRFIA denote the mean ready-for-issue assets, TCOTM the total completed on-time maintenance, and MIT the mean idle technicians. For a given strategy x and scenario ω, the AMMM produces the outcomes ($\mathrm{MRFIA}_{x,\omega}$, $\mathrm{TCOTM}_{x,\omega}$, $\mathrm{MIT}_{x,\omega}$). The multi-criteria approach manages trade-offs between maximising asset readiness, maximizing on-time maintenance completion, and minimising workforce idle time. We evaluate the robustness of these objectives by computing selected quantiles across all scenarios, namely the 10th percentile MRFIA, the 10th percentile TCOTM, and the 90th percentile MIT, by generating multiple scenarios from the uncertainty space. Higher values of the first two quantiles indicate better performance, while a lower value of the third indicates more efficient workforce utilisation. Together, these robustness metrics form a vector-valued objective function used to compare and rank strategies under uncertainty. The constraints of the problem—such as workforce availability, competency progression, maintenance capacity, resource failures, and acquisition schedules—are implicitly enforced through the simulation model's state-transition logic and business rules, ensuring that each strategy x respects the operational limits of the system.

Next, we discuss the performance measures. In the sense of an application in the military, under any asset acquisition policy, we need to maintain the maximum number of ready-for-issue assets throughout the planning horizon to use them when needed (Turan *et al*., 2020).

Performance measure 1: The mean ready-for-issue assets, MRFIA

$$MRFIA(x) = \frac{1}{T}\sum_{t=1}^{T}\sum_{i\in A} A_i(t|x)$$

$$A_i(t|x) = \begin{cases} 1 & \text{if asset } i \text{ is in the ready} - \text{for} - \text{issue state at time } t \text{ under strategy } x \\ 0 & \text{otherwise} \end{cases}$$

Higher recruitment may maintain a higher ready-for-issue assets with the price of more technicians being idle. So, we set our second performance measure to the average idle technicians and is mathematically given below.

Performance measure 2: The mean idle technicians, MIT

$$MIT(x) = \frac{1}{T}\sum_{t=1}^{T}\sum_{j\in J} IT_j(t|x)$$

where $IT_j(t|x)$ denotes the number of idle technicians with competency level j at time t under strategy x.

RFI assets may hide the information of whether assets being maintained are on-time or not. Thus, our third performance measure is total completed on-time maintenance, which capture the aspects of scheduled maintenance, workforce, and resources, and is mathematically given below.

Performance measure 3: The total completed on-time maintenance, TCOTM

$$TCOTM(x) = \frac{1}{T}\sum_{t=1}^{T}\sum_{i\in A} COTM_i(t|x)$$

where $COTM_i(t|x)$ denotes the number of maintenance events completed on-time for asset i in time period t under strategy x.

4. Methods

This section describes the methodological framework used to evaluate predefined maintenance–workforce strategies under deep uncertainty. The approach integrates the AMMM with exploratory modelling and analysis to assess how each strategy performs across a wide set of plausible future conditions. Section 4.1 outlines the robust decision-making (RDM) process, including decision framing, scenario generation, vulnerability analysis, and robustness assessment. Subsequent subsections describe how scenarios are generated using Latin Hypercube sampling, how performance metrics are computed across experiments, and how scenario discovery methods are applied to identify the uncertain parameters most responsible for strategy vulnerability.

4.1 Robust decision-making framework

The notion of utilizing diverse future perspectives to examine alternative strategies is not confined to robust decision frameworks. It is also fundamental to conventional scenario planning techniques. These methods operate on the premise that a limited set of narratives depicting various potential futures can aid planners in better preparing for unexpected developments (Schwartz, 2012). Nonetheless, traditional scenario methods grapple with two crucial questions: (1) which future scenarios to emphasize? and (2) how to translate these scenarios into actionable decisions?

Robust decision making operates iteratively and serves three key purposes: (1) it aids in the identification of potentially robust strategies, (2) it characterizes vulnerabilities inherent in these strategies, and (3) it evaluates trade-offs among them (Lempert, 2003). Within the vulnerability analysis, robust decision-making employs a systematic, computer-assisted technique called *scenario discovery* to select a few combinations of uncertain parameters and their critical ranges. These combinations effectively encapsulate future scenarios where a given strategy performs poorly. The iterative process for our specific case is briefly outlined below.

Step 1 entails the decision framing process, wherein stakeholders and decision-makers lay the foundation for analysis. In this cooperative effort, participants frame the problem by (1) specifying critical performance measures or outcomes that an organization will utilize to assess the success of alternative strategies, (2) articulating a set of such strategies with the potential to accomplish these predefined outcomes, (3) identifying pivotal uncertainties concerning the future that could significantly influence the performance of these strategies, and (4) conducting computational experiments (Lempert, 2003; Weaver *et al*., 2013).

In our case, the strategies pertain to different implementations of periodic maintenance scheduling and workforce planning. The performance metrics encompass mean ready-for-issue assets, average idle workforce, and total on-time maintenance completions. The uncertainties revolve around workforce separation across technician competency levels, failure rates, and repair times for durable maintenance resources. Section 3.2 defines the performance measures, parameters and variables. Figure 2 shows the logical structure of the exploratory modelling and analysis workbench used to configure strategies, uncertainties, and outcomes for the model. This structure streamlines the execution of computational experiments, where each experiment corresponds to a model run within a specific scenario.

Step 2 conducts computational experiments evaluating each strategy against the same set of scenarios or plausible futures. A scenario is a point within a large multi-dimensional space of possible futures or simply the uncertainty space (Bowden *et al*., 2015). Scenarios are generated by sampling the uncertainty space. We use Latin Hypercube sampling, with outcomes stored for the rest of the analysis.

Step 3 is a vulnerability analysis, focussing on identifying critical factors and their critical ranges where vulnerabilities emerge characterized by the definition of success. This is normally a threshold value which aids to classify computational experiments being decision

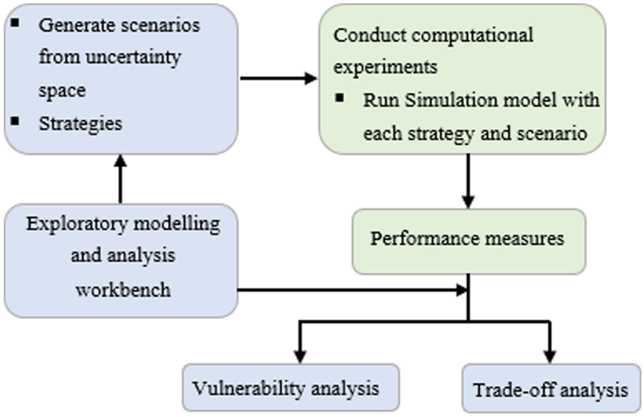

Figure 2. Integrated view of the exploratory modelling and analysis workbench and the AMMM

relevant or not. Robust decision-making conducts a series of computational experiments by executing simulation runs with many plausible future conditions which are then used to identify the decision-relevant scenarios.

To identify these critical uncertainty factors and their critical ranges, scenario discovery is used. Scenario discovery is the analytical core of robust decision-making (Kwakkel *et al.*, 2016; Bryant and Lempert, 2010). It identifies one or more sub-spaces within the uncertainty space associated with a model that is decision-relevant or vulnerable (Kwakkel, 2019). This process entails identifying the critical uncertainties, including precise variable ranges in which vulnerabilities manifest (Hadka *et al.*, 2015). Additionally, scenario discovery sheds light on which uncertain parameters play a lesser role in describing vulnerable outcomes. The most popular algorithms in the literature for scenario discovery are patient rule induction (Friedman and Fisher, 1999) and classification and regression trees (Breiman *et al.*, 2017). Patient rule induction, a statistical machine learning technique, is the most commonly and frequently used algorithm in scenario discovery (Kwakkel and Cunningham, 2016).

We use patient rule induction, leveraging exploratory modelling and analysis for scenario discovery. It identifies the scenario boxes (in the form of hyperrectangular scenario boxes), which maximize the coverage and the density and minimize the interpretability (Steinmann *et al.*, 2020). Coverage is the fraction of decision-relevant inputs out of all the decision-relevant inputs contained within the box (i.e. if all decision-relevant inputs are inside the box, then coverage is 1). Density is the fraction of inputs within the box that is decision-relevant (i.e. if all the points in the box are decision-relevant inputs, then the density is 1). Interpretability refers to the count of parameter space dimensions along which the box's size is constrained. Ultimately, patient rule induction identifies a trade-off curve, which comprises a collection of boxes situated on a Pareto optimal surface characterized by coverage, density, and interpretability.

Step 4 conducts the trade-off analysis where it finds robust strategies according to robustness metrics that are derived from outcomes. The robustness metrics indicate the decision-maker's robustness requirement. We have a series of outcomes for a strategy as each strategy is evaluated against many scenarios. So, we may have several choices for robustness metrics depending on the requirement as pointed out by McPhail *et al.* (2018). Our aim is to identify the most robust strategy among a given set of strategies.

Step 5 redesigns strategies to improve their robustness by considering new futures and results from Steps 3 and 4. In essence, vulnerability analysis supports redesigning of strategies.

5. Analysis of the results

This section presents the results of the computational experiments used to evaluate the predefined maintenance–workforce strategies. The asset acquisition policy underlying all experiments specifies an initial stock of 12 ready-for-issue assets and 4 assets requiring maintenance. Additional acquisitions occur every two years on January 1st, with quantities of 2, 2, 2, and 1 assets over the 10-year planning horizon. Using this acquisition schedule as the baseline context, we compare the performance of all nine strategies across key outcome measures to identify the most robust option. Section 5.1 reports the strategy-level comparison using robustness metrics derived from the simulation ensemble, followed by an examination of the distribution of performance outcomes to explain why the preferred strategy performs consistently well. The section concludes with the vulnerability analysis for the preferred strategy, highlighting the uncertain parameters and critical ranges that most strongly influence system under performance.

5.1 Analysis of nine different strategies

By following *Step 1* in Section 4.1, we selected nine different strategies containing different realizations of workforce and periodic maintenance plans, which are shown in Table 1. We identified the uncertainties for the analysis, which are tabulated in Table 2. We used the performance measures discussed in Section 3.2 to asses strategies. Next, we tested each strategy against the same set of 1,000 scenarios following *Step 1* in Section 4.1. That is, for each strategy, we run the simulation model 1,000 times with each scenario. The simulation model runs with an hourly time step over a 10-year planning horizon. Scenarios are generated using Latin Hypercube sampling from the uncertainty space, which is shown in Table 2.

Table 1. Nine strategies comprising different combinations of workforce plans and periodic maintenance plans

Strategy	Composition
Strategy 1	$[\boldsymbol{W}^h = [\alpha^M = 20, \alpha^V = 16, \alpha^I = 10, \alpha^C = 8], \boldsymbol{M}^h = [num\_e_i = 4, num\_a_i = 4]]$
Strategy 2	$[\boldsymbol{W}^h = [\alpha^M = 20, \alpha^V = 16, \alpha^I = 10, \alpha^C = 8], \boldsymbol{M}^m = [num\_e_i = 2, num\_a_i = 2]]$
Strategy 3	$[\boldsymbol{W}^h = [\alpha^M = 20, \alpha^V = 16, \alpha^I = 10, \alpha^C = 8], \boldsymbol{M}^l = [num\_e_i = 1, num\_a_i = 1]]$
Strategy 4	$[\boldsymbol{W}^m = [\alpha^M = 10, \alpha^V = 8, \alpha^I = 5, \alpha^C = 4], \boldsymbol{M}^h = [num\_e_i = 4, num\_a_i = 4]]$
Strategy 5	$[\boldsymbol{W}^m = [\alpha^M = 10, \alpha^V = 8, \alpha^I = 5, \alpha^C = 4], \boldsymbol{M}^m = [num\_e_i = 2, num\_a_i = 2]]$
Strategy 6	$[\boldsymbol{W}^m = [\alpha^M = 10, \alpha^V = 8, \alpha^I = 5, \alpha^C = 4], \boldsymbol{M}^l = [num\_e_i = 1, num\_a_i = 1]]$
Strategy 7	$[\boldsymbol{W}^l = [\alpha^M = 4, \alpha^V = 3, \alpha^l = 2, \alpha^C = 1], \boldsymbol{M}^h = [num\_e_i = 4, num\_a_i = 4]]$
Strategy 8	$[\boldsymbol{W}^l = [\alpha^M = 4, \alpha^V = 3, \alpha^l = 2, \alpha^C = 1], \boldsymbol{M}^m = [num\_e_i = 2, num\_a_i = 2]]$
Strategy 9	$[\boldsymbol{W}^l = [\alpha^M = 4, \alpha^V = 3, \alpha^l = 2, \alpha^C = 1], \boldsymbol{M}^l = [num\_e_i = 1, num\_a_i = 1]]$

Table 2. Uncertainty space

Uncertain model parameter	Range	Integer/ real
s^M	[0, 0.02]	Real
s^V	[0, 0.02]	Real
s^I	[0, 0.02]	Real
s^C	[0, 0.02]	Real
s^{Ma}	[0, 0.02]	Real
fr_1	[0.05, 0.15]	Real
rt_1	[15, 22]	Integer
fr_2	[0.05, 0.25]	Real
rt_2	[12, 18]	Integer

We score each ensemble outcome (three performance measures) and corresponding input (scenario values) to identify the relative influence of each input on the outcomes. This was accomplished using the "Extra Trees Algorithm" from machine learning (Geurts *et al*., 2006). The results of this scoring (Figure 3) indicate that the outcomes are heavily influenced by the strategies. Thus, it is important to select the right strategy to achieve better performance.

Continuing with Step 4 (Section 4.1), we prescribe the robustness requirement for the trade-off analysis. We use descriptive statistics (10th percentile worst case scenarios) when defining robustness metrics (McPhail *et al*., 2018). Note that we have results from 1,000 computational experiments for each strategy, i.e. we have an array of 1,000 values for each outcome.

(1) The 10th percentile mean ready-for-issue assets (10MRFIA). Higher values indicate better robustness, as we aim to maintain a higher level of ready-for-issue assets across all scenarios.

(2) The 10th percentile total completed on-time maintenance (10TCOTM). Higher values indicate better robustness, as we aim to complete maintenance tasks on time across all scenarios.

(3) The 90th percentile mean idle technicians (90MIT). Lower values indicate better robustness, as we aim to maintain a minimal level of idle technicians across all scenarios. This also implicitly reflects cost effectiveness.

In this study, equal weights are adopted as a neutral starting point to enable a transparent comparison of strategies without privileging any single outcome in advance. This approach is commonly used in exploratory analyses when decision-makers have not specified formal preferences or when multiple outcomes are regarded as broadly important. Equal weighting therefore serves as a baseline for examining the relative behaviour of strategies under deep uncertainty rather than as a prescriptive statement about organisational priorities.

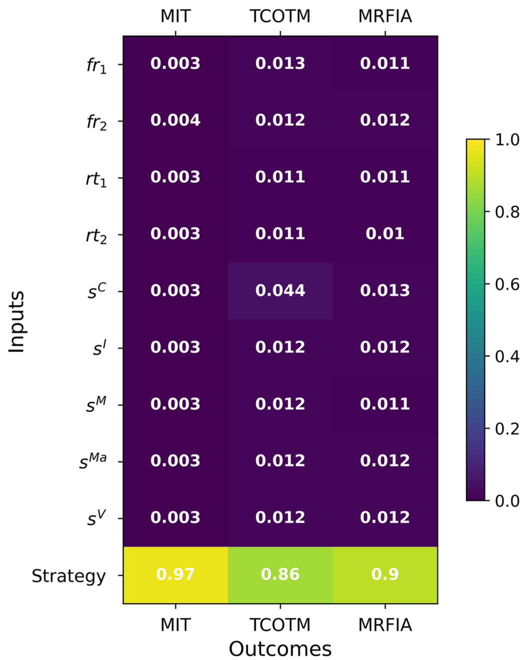

Figure 3. Feature scoring results

With the aid of this robustness analysis, we evaluate the performance of the nine strategies across the simulated scenarios. Table 3 reports the robustness values for each strategy, while Figure 4 presents these results using a parallel coordinate plot. Figure 4 shows that Strategy 7, composed of a low workforce ($\mathbf{W}^l$) and a high maintenance plan ($\mathbf{M}^h$), performs best across all three robustness metrics. This is indicated by the straight line for Strategy 7 appearing at the top of the plot. As Strategy 7 outperforms all other strategies across the three robustness metrics, it is identified as the preferred strategy. This result is not immediately evident without the simulation analysis. For example, one might expect higher values of the 10MRFIA and 10TCOTM metrics for the strategy combining a high workforce ($\mathbf{W}^h$) with a high maintenance plan ($\mathbf{M}^h$) (i.e. Strategy 1). This finding highlights the value of the simulation model in capturing the complex trade-offs between maintenance planning and workforce decisions under deep system uncertainty.

The parallel coordinate plot (Figure 4) also illustrates a trade-off apparent between the 90MIT and the 10TCOTM, as evidenced by the intersecting lines along these two dimensions. Additionally, a correlation is apparent between the 10TCOTM and the 10MRFIA, since the lines do not intersect across these dimensions.

Next, we analyse the results from the 1,000 computational experiments conducted for each strategy to gain further insight into why Strategy 7 is preferred. Figure 5(a) and (b) present box plots showing the distribution of each performance measure for all strategies, using the standard five-number summary (minimum, first quartile, median, third quartile, and maximum). Several patterns emerge from these results. First, we observe that all three maintenance schedules produce more idle technicians when annual recruitment is high, with idle time decreasing as recruitment moves from medium to low; thus, higher annual recruitment is not desirable. Second, high maintenance frequency results in fewer idle technicians and higher completed on-time maintenance and ready-for-issue assets across all workforce plans; however, it also introduces greater variability in these performance measures due to reduced availability of mentor technicians when separation rates are high. This lack of mentors increases idle technicians and decreases completed on-time maintenance, which is reflected in the wider variation of ready-for-issue assets. When robustness is defined in terms of the range of performance values, medium or low maintenance schedules are therefore preferred because they produce more stable outcomes. Third, when comparing the high-maintenance schedule across workforce plans, the low-workforce plan performs best under the robustness metrics discussed earlier, and this preference also holds for the other two maintenance schedules. Fourth, for the high-maintenance schedule, the mean number of idle technicians decreases from the high-workforce plan to the low-workforce plan, while the total completed on-time maintenance and the mean ready-for-issue assets remain comparable or even improve in some scenarios. This again reflects the shortage of mentor technicians in scenarios with high recruitment and high separation rates, which increases idle technicians and reduces maintenance completion and ready-for-issue assets. As recruitment decreases, the effect of mentor shortages becomes less severe, as shown by the narrowing ranges. Finally,

Table 3. Robustness metrics (90MIT, 10TCOTM, 10MRFIA) for nine strategies

Strategy	90MIT	10TCOTM	10MRFIA
Strategy 1	246.77	63.0	4.33
Strategy 2	247.93	18.0	1.11
Strategy 3	248.01	5.0	0.43
Strategy 4	126.82	83.0	5.33
Strategy 5	128.55	24.0	1.44
Strategy 6	128.72	7.0	0.49
Strategy 7	51.41	122.0	5.33
Strategy 8	53.63	31.0	1.44
Strategy 9	53.88	10.0	0.56

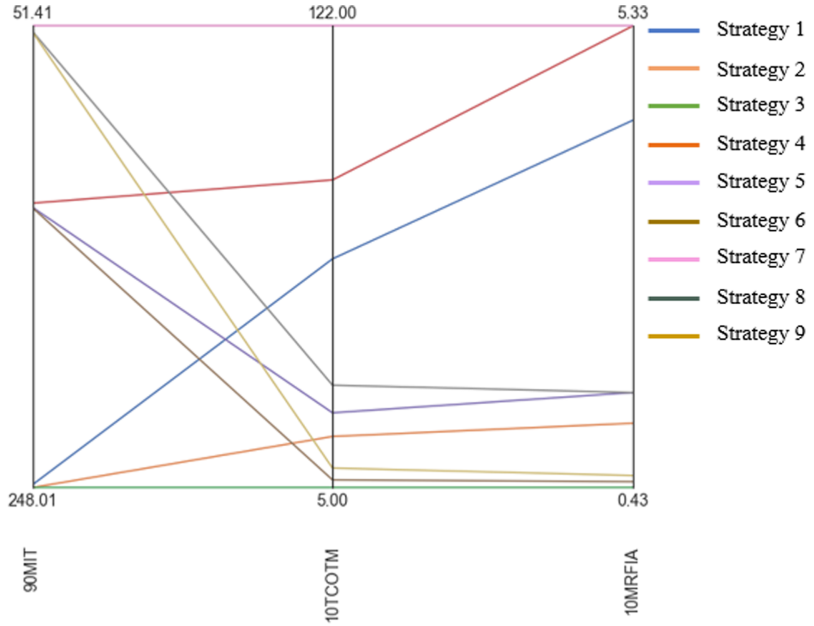

Figure 4. Parallel coordinate plot for nine strategies. The 90MIT axis is inverted for readability

although low recruitment under the high-maintenance schedule produces a slightly lower maximum value for total completed on-time maintenance than the high- or medium-workforce plans (with maxima of 134, 135, and 127 respectively), it maintains consistently high performance across all scenarios. This is supported by the stable mean ready-for-issue asset levels, indicating that outstanding maintenance is eventually completed under all scenarios.

This deeper consideration of workload due to mentoring reveals additional value from the robustness of Strategy 7.

5.2 Vulnerability analysis of the preferred robust strategy

Next, we conduct a vulnerability analysis of this strategy following *Step 3* in Section 4.1. In particular, we seek to generate decision-making insights regarding which uncertain model parameters are most critical and the ranges of these parameters that explain poor performance in scenarios characterized by the definition of success. Figure 6 shows the results from 1,000 computational experiments for the robust strategy. This parallel coordinate plot shows both scenarios and performance measures. We observe that the MRFIA is identical in all the scenarios (Figure 6, and the bottom-right plot of Figure 5(b) and, hence, we exclude it from the vulnerability analysis. Instead, the analysis investigates why the strategy performs poorly (or differently) across alternative scenarios and what causes these behaviours. Therefore, the remainder of the analysis focuses on the TCOTM and the MIT.

Vulnerability analysis helps identify the critical uncertain model parameters and the ranges in which the system is most likely to underperform if the strategy is implemented. These critical parameters represent the factors that must be monitored to avoid system failure, while the critical ranges indicate when adjustment to the strategy might be required.

5.2.1 What causes lower on-time completion of maintenance?. In this analysis, we examine why total completed on-time maintenance falls below an acceptable performance threshold in certain scenarios. Following McPhail *et al.* (2018), we define this threshold as the 75th percentile of total on-time maintenance in the baseline ensemble. Scenarios that fall below this

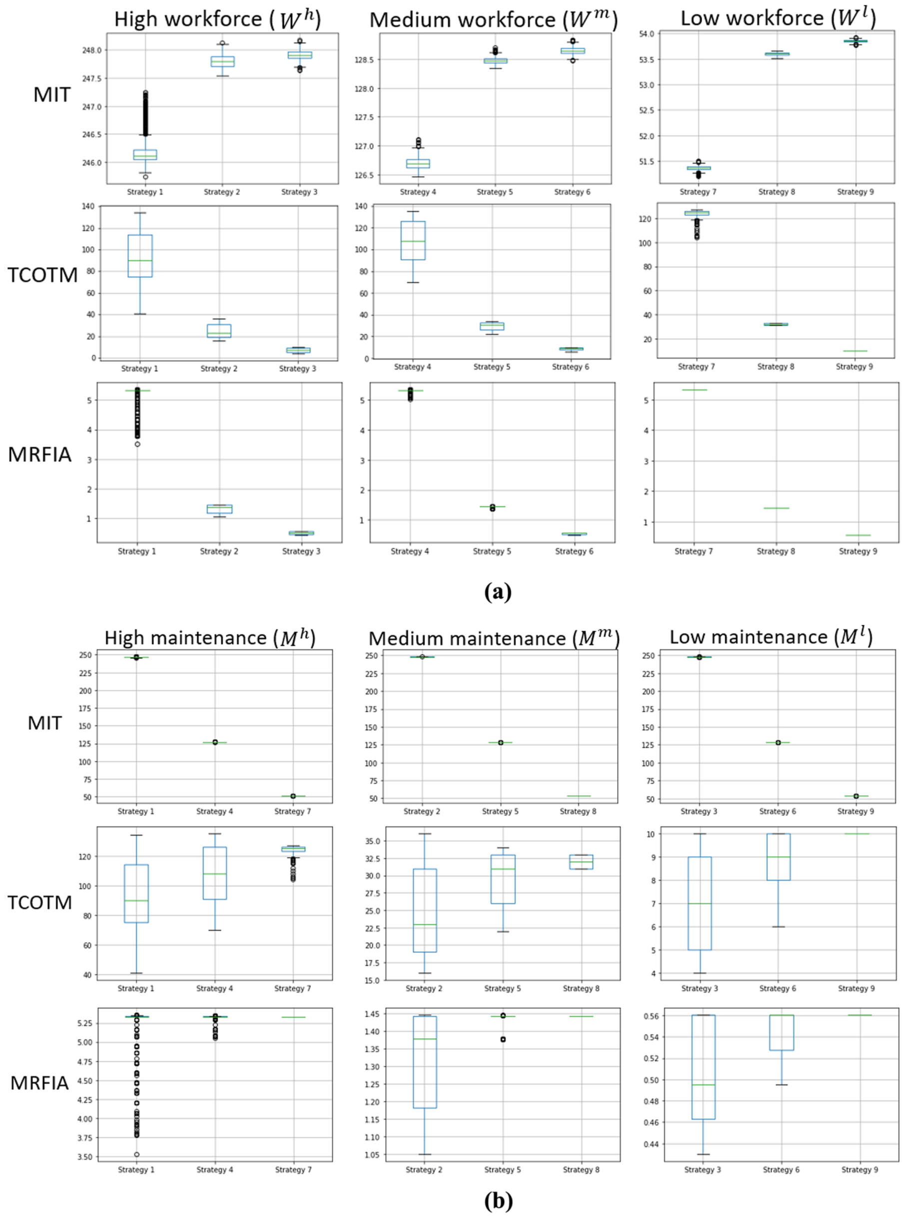

Figure 5. Distribution of performance measures across 1,000 experiments for each strategy, grouped by (a) Workforce intensity and (b) Maintenance schedule intensity

level are classified as vulnerable. Out of the 1,000 computational experiments conducted for the preferred strategy, 723 meet this definition.

Scenario discovery using patient rule induction (PRIM) yields a sequence of candidate scenario boxes, each representing combinations of uncertain factors associated with vulnerability. Figure 7(a) presents the density–coverage frontier for these boxes. We

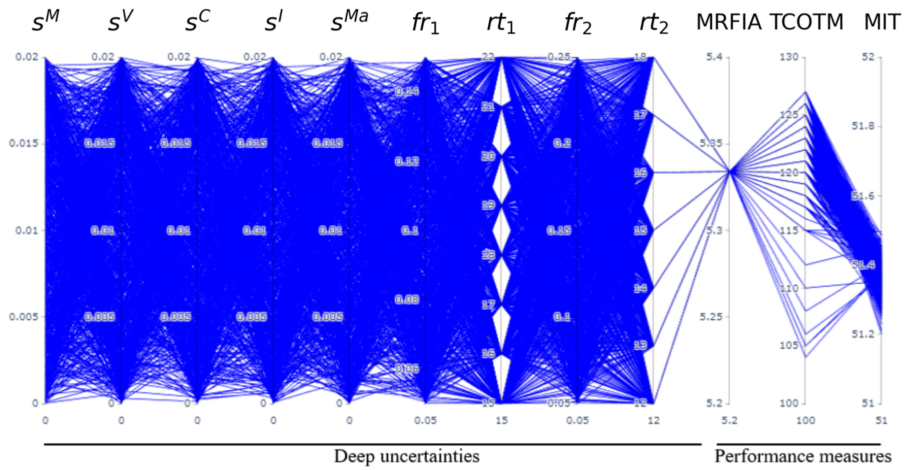

Figure 6. Parallel coordinate plot of 1,000 experiments for the preferred strategy

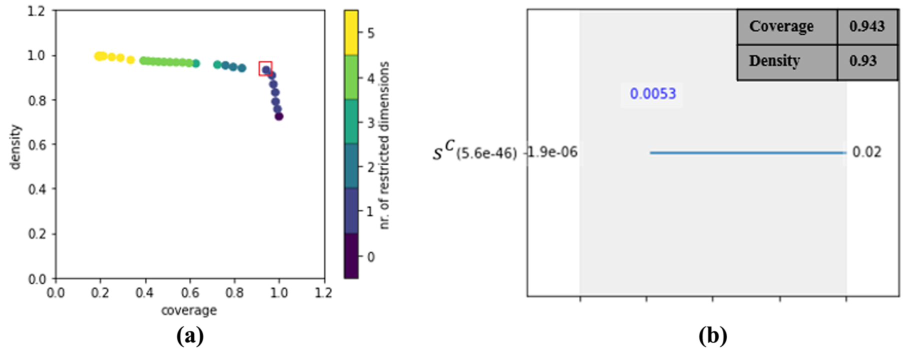

Figure 7. Scenario discovery results for low on-time maintenance completion: (a) Density–coverage trade-off, and (b) Corresponding scenario box

selected a region that balances high coverage with sufficient density while limiting the number of restricted dimensions for interpretability. This box captures slightly more than 94% of all vulnerable scenarios, and its boundaries are shown in Figure 7(b).

The analysis identifies one dominant driver of underperformance: the separation fraction of coordinators (s^C). When s^C exceeds 0.0053, approximately 93% of simulation runs fall below the threshold, indicating that on-time maintenance completion is highly sensitive to coordinator availability. This pattern aligns with the results in Figure 6, where mean ready-for-issue asset levels remain relatively stable across scenarios using the robust strategy. This suggests that higher coordinator turnover manifests primarily as delayed task completion and backlog accumulation rather than permanent loss of operational capacity.

From a decision-making perspective, s^C represents a critical uncertainty that warrants continuous monitoring. The results have two practical implications. First, robustness can be enhanced by reducing coordinator turnover or by ensuring timely replacement and training. Second, trends in s^C provide an early warning signal: as values approach the critical range identified in Figure 7(b), proactive measures—such as targeted promotion or accelerated recruitment—may be required to prevent degradation in maintenance performance.

5.2.2 What causes higher idle technicians?. Following the analysis of on-time maintenance performance, we next examine why the mean idle technicians exceeds an acceptable threshold in some scenarios. We define the vulnerability threshold as the 25th percentile of MIT in the baseline results. Scenarios with idle levels above this value are classified as vulnerable. Under this definition, 750 out of 1,000 computational experiments are vulnerable.

Scenario discovery using PRIM generates a series of candidate boxes describing combinations of uncertain parameters associated with elevated idle technician levels. Figure 8(a) presents the density–coverage frontier for these boxes. We selected a box that provides a strong balance between density and coverage while maintaining interpretability by limiting the number of restricted dimensions. The chosen box captures approximately 80% of all vulnerable scenarios, with boundaries shown in Figure 8(b). The analysis identifies a single dominant driver of high idle technician levels: the separation fraction of verifiers (s^V). When $s^V < 0.015$, roughly 82% of all experiments exceed the idle workforce threshold. This indicates that idle workforce levels under the robust strategy (Strategy 7) are particularly sensitive to the availability of verifiers. As also reflected in Figure 6, verifier separations influence idle levels without materially affecting total on-time maintenance completion, implying a structural imbalance where the system contains more verifiers than needed under most scenarios.

From a decision-making perspective, s^V emerges as a critical uncertainty that should be monitored for long-term workforce sustainability. Two implications follow. First, the workforce plan could be modestly refined by reducing verifier recruitment rates, thereby reducing idle technician levels without compromising maintenance throughput. Second, because higher workforce pressure can increase turnover among specialist roles, trends in s^V may serve as an early signal of imbalance.

Although minimising idle technician time is desirable from an efficiency perspective, it is important to recognise that very low idle time may place additional pressure on key personnel such as coordinators, inspectors, and verifiers. Empirical evidence from both military and civilian maintenance environments shows that high operational tempo, especially when combined with limited recruitment, can contribute to increased turnover among experienced staff (Olsen, 2008). This interaction suggests that workforce plans aimed at reducing idle time should also consider the long term sustainability of the workforce, as higher attrition in these critical roles may undermine the effectiveness of otherwise robust maintenance strategies.

Next, we assess the relative influence of the uncertain parameters on the outcomes using feature scoring, excluding MRFIA since it does not vary across scenarios. As shown in Figure 9, the separation fraction of coordinators (s^C) is the dominant driver of total on-time

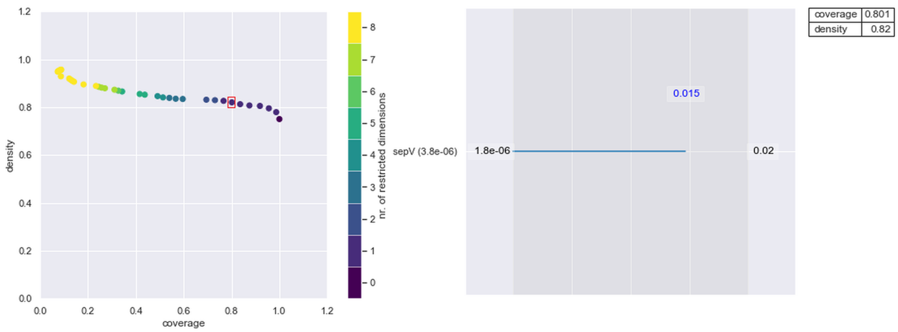

Figure 8. Scenario discovery results for high idle technician levels: (a) Density–coverage trade-off, and (b) Corresponding scenario box

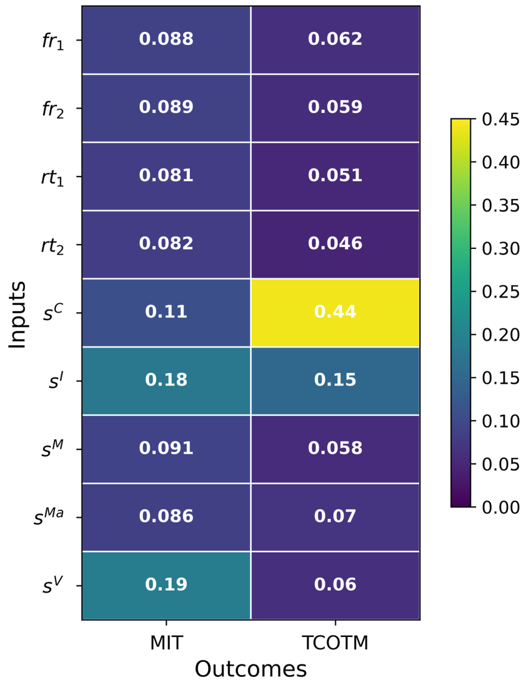

Figure 9. Feature scoring results showing the relative influence of uncertain parameters on outcomes for the preferred strategy

maintenance, followed by inspectors (s^I), consistent with the scenario-discovery results. For idle technicians, the separation fraction of verifiers (s^V) has the strongest influence, again followed by s^I and s^C. Other parameters have comparatively minor effects. Overall, coordinator, verifier, and inspector separation rates emerge as the most critical uncertainties and should be monitored closely to prevent degradation in system performance.

6. Conclusion

Several key findings and suggestions emerge from this study. First, in the early stages of planning, it is essential to define scenarios that represent plausible future conditions, as this supports the development of robust and resilient strategies—especially when historical data are unavailable or limited, making parameter estimation challenging. Second, Pareto-optimal solutions identified through simulation–optimization under best-estimate assumptions may be misleading (Lempert, 2019), and given that skilled workforce availability is often the main constraint in maintenance environments, strategies produced solely by optimization methods may not be workable in practice. Exploratory modelling and analysis therefore provide a valuable means to support strategic maintenance and workforce planning by enabling more informed and resilient decisions. Third, the findings highlight the importance of balancing maintenance scheduling with workforce planning to enhance workforce performance, asset readiness, and on-time maintenance completion. Finally, the results reinforce the need for a well-aligned workforce plan to consistently achieve the desired performance outcomes across a broad range of future conditions.

The findings demonstrate that integrating a detailed maintenance workforce simulation model with exploratory modelling and analysis provides a practical and rigorous basis for planning under deep uncertainty. By highlighting how maintenance schedules and workforce recruitment interact in a constrained skilled-workforce environment, and by identifying the conditions under which even robust strategies may become vulnerable, the study supports more informed and resilient decision-making for high-value defence assets.

While the framework provides meaningful insights into maintenance and skilled-workforce planning under deep uncertainty, several limitations should be acknowledged. The current model focuses on a single asset type and assumes fixed acquisition timings, constant separation rates, and scenario-dependent but time-invariant failure and repair rates for maintenance resources. These simplifications enable tractable analysis but may not fully capture the evolving operational environment or shifts in workforce behaviour over long planning horizons. The model also abstracts certain aspects of technician training pathways and cross-platform skill transfer, which may be relevant in more complex real-world settings.

From a managerial perspective, the results highlight the importance of jointly considering maintenance scheduling and workforce recruitment decisions rather than treating them as separate planning tasks. The robustness analysis indicates that strategies performing well under best-estimate conditions may fail when uncertainties in separation rates, mentor availability, or resource failures materialise. Managers can therefore use the global sensitivity analysis to identify parameters requiring continuous monitoring, and employ the vulnerability analysis outputs as early-warning indicators when shifts in uncertainty may push the system toward under-performance. These insights can support defence organisations in designing more resilient maintenance and workforce plans, prioritising training investments, and anticipating conditions under which strategic adjustments may be required.

Acknowledgments

The authors used a large language model (LLM) to assist with editorial refinement and clarity improvements during the manuscript revision. All substantive ideas, analyses, and contributions are solely those of the authors.

We acknowledge that the analyses in this paper are not a reflection of the position, intent or opinions of any defense organization, but are solely the opinions of the authors of the paper.

References

Akl, A.M., El Sawah, S., Chakrabortty, R.K. and Turan, H.H. (2022), “A joint optimization of strategic workforce planning and preventive maintenance scheduling: a simulation–optimization approach”, *Reliability Engineering and System Safety*, Vol. 219, 108175, doi: 10.1016/j.ress.2021.108175.

Bowden, F.D., Pincombe, B. and Williams, P.B. (2015), “Feasible scenario spaces: a new way of measuring capability impacts”, *MODSIM2015*, pp. 836-842.

Breiman, L., Friedman, J., Olshen, R.A. and Stone, C.J. (2017), *Classification and Regression Trees*, Chapman and Hall/CRC, New York.

Bryant, B.P. and Lempert, R.J. (2010), “Thinking inside the box: a participatory, computer-assisted approach to scenario discovery”, *Technological Forecasting and Social Change*, Vol. 77 No. 1, pp. 34-49, doi: 10.1016/j.techfore.2009.08.002.

De Bruecker, P., Van den Bergh, J., Beliën, J. and Demeulemeester, E. (2015), “Workforce planning incorporating skills: state of the art”, *European Journal of Operational Research*, Vol. 243, pp. 1-16, doi: 10.1016/j.ejor.2014.10.038.

Dong, W., Liu, S. and Du, Y. (2019), “Optimal periodic maintenance policies for a parallel redundant system with component dependencies”, *Computers and Industrial Engineering*, Vol. 138, 106133, doi: 10.1016/j.cie.2019.106133.

Friedman, J.H. and Fisher, N.I. (1999), "Bump hunting in high-dimensional data", *Statistics and Computing*, Vol. 9 No. 2, pp. 123-143, doi: 10.1023/a:1008894516817.

Fu, Y. and Wang, J. (2022), "Optimum periodic maintenance policy of repairable multi-component system with component reallocation and system overhaul", *Reliability Engineering and System Safety*, Vol. 219, 108224, doi: 10.1016/j.ress.2021.108224.

Geurts, P., Ernst, D. and Wehenkel, L. (2006), "Extremely randomized trees", *Machine Learning*, Vol. 63 No. 1, pp. 3-42, doi: 10.1007/s10994-006-6226-1.

Hadka, D., Herman, J., Reed, P. and Keller, K. (2015), "An open source framework for many-objective robust decision making", *Environmental Modelling and Software*, Vol. 74, pp. 114-129, doi: 10.1016/j.envsoft.2015.07.014.

Hall, J. (2007), "Probabilistic climate scenarios may misrepresent uncertainty and lead to bad adaptation decisions", *Hydrological Processes*, Vol. 21 No. 8, pp. 1127-1129.

Horn, M.E., Elgindy, T. and Gomez-Iglesias, A. (2016), "Strategic workforce planning for the Australian defence force", *Journal of the Operational Research Society*, Vol. 67 No. 4, pp. 664-675, doi: 10.1057/jors.2015.85.

Jafar-Zanjani, H., Zandieh, M. and Sharifi, M. (2022), "Robust and resilient joint periodic maintenance planning and scheduling in a multi-factory network under uncertainty: a case study", *Reliability Engineering and System Safety*, Vol. 217, 108113, doi: 10.1016/j.ress.2021.108113.

Kafiabad, S.T., Zanjani, M.K. and Nourelfath, M. (2020), "Integrated planning of operations and on-job training in maintenance logistics networks", *Reliability Engineering and System Safety*, Vol. 199, 106922, doi: 10.1016/j.ress.2020.106922.

Kwakkel, J.H. (2019), "A generalized many-objective optimization approach for scenario discovery", *Futures and Foresight Science*, Vol. 1 No. 2, e8, doi: 10.1002/ffo2.8.

Kwakkel, J.H. and Cunningham, S.C. (2016), "Improving scenario discovery by bagging random boxes", *Technological Forecasting and Social Change*, Vol. 111, pp. 124-134, doi: 10.1016/j.techfore.2016.06.014.

Kwakkel, J.H., Haasnoot, M. and Walker, W.E. (2016), "Comparing robust decision-making and dynamic adaptive policy pathways for model-based decision support under deep uncertainty", *Environmental Modelling and Software*, Vol. 86, pp. 168-183, doi: 10.1016/j.envsoft.2016.09.017.

Lempert, R.J. (2003), *Shaping the Next One Hundred Years: New Methods for Quantitative, Long-Term Policy Analysis*, Rand Corporation, Santa Monica, CA.

Lempert, R.J. (2019), "Robust decision making (RDM)", in *Decision Making under Deep Uncertainty*, Springer, Cham, pp. 23-51.

Lempert, R.J., Groves, D.G., Popper, S.W. and Bankes, S.C. (2006), "A general, analytic method for generating robust strategies and narrative scenarios", *Management Science*, Vol. 52 No. 4, pp. 514-528, doi: 10.1287/mnsc.1050.0472.

Liao, C.-J. and Chen, W.-J. (2003), "Single-machine scheduling with periodic maintenance and nonresumable jobs", *Computers and Operations Research*, Vol. 30 No. 9, pp. 1335-1347, doi: 10.1016/s0305-0548(02)00074-6.

Mansour, M.A.A.-F. (2011), "Solving the periodic maintenance scheduling problem via genetic algorithm to balance workforce levels and maintenance cost", *American Journal of Engineering and Applied Sciences*, Vol. 4 No. 2, pp. 223-234, doi: 10.3844/ajeassp.2011.223.234.

McPhail, C., Maier, H.R., Kwakkel, J.H., Giuliani, M., Castelletti, A. and Westra, S. (2018), "Robustness metrics: how are they calculated, when should they be used and why do they give different results?", *Earth's Future*, Vol. 6 No. 2, pp. 169-191, doi: 10.1002/2017ef000649.

Morgan, M.G., Henrion, M. and Small, M. (1990), *Uncertainty: A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*, Cambridge University Press, Cambridge.

Olsen, N.P. (2008), “The impact of operations tempo (OPTEMPO) on intentions to depart the military. Does the increase of OPTEMPO cause action”, Technical Report.

Perez-Gonzalez, P., Fernandez-Viagas, V. and Framinan, J.M. (2020), “Permutation flowshop scheduling with periodic maintenance and makespan objective”, *Computers and Industrial Engineering*, Vol. 143, 106369, doi: 10.1016/j.cie.2020.106369.

Safaei, N., Banjevic, D. and Jardine, A.K. (2011), “Workforce-constrained maintenance scheduling for military aircraft fleet: a case study”, *Annals of Operations Research*, Vol. 186 No. 1, pp. 295-316, doi: 10.1007/s10479-011-0885-4.

Sbihi, M. and Varnier, C. (2008), “Single-machine scheduling with periodic and flexible periodic maintenance to minimize maximum tardiness”, *Computers and Industrial Engineering*, Vol. 55 No. 4, pp. 830-840, doi: 10.1016/j.cie.2008.03.005.

Scholte, M., Marchau, V.A., Kwakkel, J.H., Klijn, C.J., Rovers, M.M. and Grutters, J.P. (2023), “Dealing with uncertainty in early health technology assessment: an exploration of methods for decision making under deep uncertainty”, *Value in Health*, Vol. 26 No. 5, pp. 694-703, doi: 10.1016/j.jval.2022.08.012.

Schwartz, P. (2012), *The Art of the Long View: Planning for the Future in an Uncertain World*, Currency, New York.

Shahmoradi-Moghadam, H., Safaei, N. and Sadjadi, S.J. (2021), “Robust maintenance scheduling of aircraft fleet: a hybrid simulation-optimization approach”, *IEEE Access*, Vol. 9, pp. 17854-17865, doi: 10.1109/access.2021.3053714.

Shen, J. and Zhu, Y. (2019), “A parallel-machine scheduling problem with periodic maintenance under uncertainty”, *Journal of Ambient Intelligence and Humanized Computing*, Vol. 10 No. 8, pp. 3171-3179, doi: 10.1007/s12652-018-1032-8.

Stanton, M.C.B. and Roelich, K. (2021), “Decision making under deep uncertainties: a review of the applicability of methods in practice”, *Technological Forecasting and Social Change*, Vol. 171, 120939, doi: 10.1016/j.techfore.2021.120939.

Steinmann, P., Auping, W.L. and Kwakkel, J.H. (2020), “Behavior-based scenario discovery using time series clustering”, *Technological Forecasting and Social Change*, Vol. 156, 120052, doi: 10.1016/j.techfore.2020.120052.

Turan, H.H., Elsawah, S. and Ryan, M.J. (2020), “A long-term fleet renewal problem under uncertainty: a simulation-based optimization approach”, *Expert Systems with Applications*, Vol. 145, 113158, doi: 10.1016/j.eswa.2019.113158.

Turan, H.H., Jalalvand, F., Kahagalage, S. and El Sawah, S. (2021a), “Integrating decision maker preferences to a risk-averse multi-objective simulation-based optimization for a military workforce planning, asset management and fleet management problem”, *Computers and Industrial Engineering*, Vol. 162, 107752, doi: 10.1016/j.cie.2021.107752.

Turan, H.H., Kahagalage, S.D., Jalalvand, F. and El Sawah, S. (2021b), “A multi-objective simulation–optimization for a joint problem of strategic facility location, workforce planning, and capacity allocation: a case study in the royal Australian navy”, *Expert Systems with Applications*, Vol. 186, 115751, doi: 10.1016/j.eswa.2021.115751.

Walker, W.E., Marchau, V.A.W.J. and Swanson, D. (2010), “Addressing deep uncertainty using adaptive policies: introduction to section 2”, *Technological Forecasting and Social Change*, Vol. 77 No. 6, pp. 917-923, doi: 10.1016/j.techfore.2010.04.004.

Weaver, C.P., Lempert, R.J., Brown, C., Hall, J.A., Revell, D. and Sarewitz, D. (2013), “Improving the contribution of climate model information to decision making: the value and demands of robust decision frameworks”, *Wiley Interdisciplinary Reviews: Climate Change*, Vol. 4 No. 1, pp. 39-60, doi: 10.1002/wcc.202.

Withers, K.R., Turan, H.H., Kahagalage, S. and El Sawah, S. (2021), “The closed loop military maintenance workforce: a simulation-based optimisation approach”, *24th International Congress on Modelling and Simulation*, Sydney, 5-10 December, pp. 785-791.

Yazdani, R., Mirmozaffari, M., Shadkam, E. and Khalili, S.M. (2023), "A lion optimisation algorithm for a two-agent single-machine scheduling with periodic maintenance to minimise the sum of maximum earliness and tardiness", *International Journal of Industrial and Systems Engineering*, Vol. 44 No. 4, pp. 515-531, doi: 10.1504/ijise.2023.132730.

Zhang, Q., Liu, Y., Xiahou, T. and Huang, H.-Z. (2023), "A heuristic maintenance scheduling framework for a military aircraft fleet under limited maintenance capacities", *Reliability Engineering and System Safety*, Vol. 235, 109239, doi: 10.1016/j.ress.2023.109239.

Corresponding author
Hasan Hüseyin Turan can be contacted at: h.turan@unsw.edu.au