Proceedings of the 11th World Congress on New Technologies (NewTech'25)
Paris, France - August, 2025
Paper No. ICCEIA 163
DOI: 10.11159/icceia25.163

Optimizing Maintenance Strategies for Transportation Networks: Integrating Risk Attitudes in Decision Making

Maria Habib[1], Rafic Faddoul[1], Wassim Raphael[1]
[1] Saint Joseph University of Beirut, Faculty of Engineering ESIB, Lebanon
Beirut, Lebanon
maria.habib2@usj.edu.lb; rafic.faddoul@usj.edu.lb; wassim.raphael@usj.edu.lb

Abstract **-** Optimizing maintenance strategies is essential for enhancing the safety, reliability, and economic efficiency of transportation infrastructure, particularly roads and bridges. Conventional optimization models typically aim to minimize costs while maximizing reliability but often overlook the varied risk attitudes of infrastructure managers and decision-makers. This limitation can lead to theoretically optimal strategies that fail to reflect practical stakeholder preferences under uncertainty. This study proposes a framework that integrates risk attitudes directly into the optimization of maintenance strategies. The approach extends a value iteration algorithm by incorporating a Constant Relative Risk Aversion (CRRA) utility function, enabling the modeling of risk-averse, risk-neutral, and risk-seeking behaviors. The framework is implemented with an actor–critic reinforcement-learning architecture, where a critic network guides a policy (actor) toward cost-optimal maintenance decisions. Unlike traditional reinforcement learning or optimization methods that assume risk neutrality, the proposed model adjusts the reward structure to reflect how decision-makers perceive the trade-offs between cost, reliability, and future uncertainty. The framework is applied to a representative road and bridge network to examine how varying risk preferences influence the selection of optimal maintenance strategies—including preventive, corrective, opportunistic, and induced maintenance. By adjusting the risk aversion parameter, the model produces different optimal strategies corresponding to diverse decision-maker profiles. To support decision-making, the study visualizes how risk attitudes affect both the selection of maintenance strategies and the associated lifecycle costs. Results highlight the significant role of risk preferences in shaping maintenance policies. This research contributes to bridging the gap between theoretical optimization techniques and the practical realities of infrastructure management by providing a flexible, risk-aware tool that enables infrastructure managers to tailor maintenance strategies to both technical requirements and organizational risk preferences.

***Keywords*:** Maintenance Optimization; Transportation Infrastructure; Risk Attitudes; CRRA Utility Function; Value Iteration; Decision Making.

1. Introduction

Maintenance of transportation infrastructure—particularly roads and bridges—is a critical function for ensuring public safety, economic continuity, and long-term system performance. The increasing age and usage intensity of transportation assets have heightened the need for maintenance strategies that are not only cost-effective but also resilient to uncertainty [1]. Traditionally, maintenance optimization models have focused on minimizing costs and maximizing reliability, often assuming that decision-makers act under risk-neutral conditions [2], [3]. However, real-world infrastructure management is more complex, as it is shaped by the risk attitudes of stakeholders, which can significantly influence maintenance priorities and investment decisions [4], [5].

Risk management has therefore become a cornerstone of modern maintenance strategy development, especially in high-stakes sectors such as transportation, where failures can lead to severe financial, environmental, or safety consequences [6]. However, most optimization models do not explicitly account for risk preferences, often assuming risk-neutral decision-making [7]. For example, a risk-averse manager might prioritize preventive maintenance to avoid catastrophic failures—even at a higher upfront cost—while a risk-seeking manager may delay interventions in pursuit of short-term savings [8].

Recent advances in reinforcement learning (RL) and Markov decision processes (MDPs) have enabled the development of adaptive maintenance optimization frameworks capable of handling complex, multi-component systems [9]. While RL methods have shown promise for optimizing maintenance under uncertainty, most applications assume risk neutrality, failing to capture the impact of risk preferences on long-term strategies. Addressing this gap requires models that can explicitly represent and incorporate risk attitudes within the optimization process.

This study proposes a risk-aware optimization framework for maintenance planning in transportation infrastructure systems, explicitly integrating decision-makers' risk attitudes into the strategy selection process. The framework is implemented through an actor–critic reinforcement-learning architecture, in which a critic network evaluates risk-adjusted state values while an actor network learns policies that minimise expected life-cycle costs under the chosen risk preference. The approach is demonstrated on a representative series-parallel network of roads and bridges within the Dutch transportation system. Results show how varying risk attitudes influence the balance of preventive, corrective, opportunistic, and induced maintenance actions, thereby shaping the overall lifecycle costs of the transportation network.

2. Methodology

2.1. Problem Formulation

The objective of this study is to optimize maintenance strategies for a transportation infrastructure system composed of roads and bridges. Each infrastructure element can exist in multiple discrete condition states, progressing over time due to deterioration or improving following maintenance interventions. The decision-maker must select, at each decision epoch, a maintenance action for each component, balancing maintenance costs, user costs (associated with reduced reliability or downtime), and risk preferences.

Each element's state reflects its condition (e.g., from "as new" to "failed"). The system state is a tuple of the condition states of all components. Each component has a discrete set of possible maintenance actions, including "do nothing" (deterioration proceeds) and various maintenance treatments (each with defined cost, effect on state transition probabilities, and downtime). State transitions follow a Markov process, for each component and action, a transition matrix defines the probabilities of moving from one condition state to another. The reward (negative cost) function accounts for direct maintenance costs and user costs related to downtime. These costs are adjusted to reflect the decision-maker's risk preferences using a utility function.

2.2. Risk Preferences Integration: CRRA Utility

To account for the varying risk preferences of infrastructure managers, this study incorporates the Constant Relative Risk Aversion (CRRA) utility function into the reward structure of the optimization model. The CRRA utility is a widely used approach in economics and decision theory for capturing how individuals or organizations perceive the trade-off between cost and risk. The utility function is defined as:

$$U(C) = \begin{cases} \dfrac{C^{1-\gamma}}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln(C) & \text{if } \gamma = 1 \end{cases} \tag{1}$$

Where C represents the total cost, comprising both maintenance expenditures and user costs associated with reduced reliability or downtime. The parameter γ denotes the coefficient of relative risk aversion: when $\gamma > 0$, the decision-maker is considered risk-averse, placing a higher weight on avoiding uncertain or catastrophic costs; when $\gamma = 0$, the decision-maker is risk-neutral, evaluating all outcomes purely based on their expected cost; and when $\gamma < 0$, the decision-maker is risk-seeking, potentially favoring strategies with lower immediate costs despite higher risks [7].

The CRRA utility was selected for this study due to its ability to model a wide range of risk attitudes with a single parameter and its compatibility with dynamic decision-making frameworks like reinforcement learning and Markov decision processes. This utility function has been successfully applied in infrastructure and insurance economics to represent managerial risk preferences [8].

2.3. Actor-Critic Reinforcement Learning Framework

The optimization of maintenance strategies in this study is achieved through an actor–critic reinforcement-learning (RL) algorithm grounded in value-iteration principles, adapted to incorporate risk attitudes via the CRRA utility function. In this actor–critic setup, the *actor* (policy network) proposes maintenance actions while the *critic* (value network) evaluates their long-term utility, allowing both components to improve iteratively.

Value iteration is a dynamic programming method widely used to solve Markov Decision Processes (MDPs), where the goal is to find a policy that maximizes the expected cumulative reward (or minimizes cumulative cost) over time. In the context of infrastructure maintenance, this corresponds to identifying the sequence of maintenance actions that minimizes the total expected cost, accounting for both direct maintenance expenses and user costs arising from reduced reliability or service interruptions. Value iteration is used in this application because of its robustness and its suitability for problems with discrete state and action spaces—common characteristics of infrastructure maintenance scenarios [10], [11].

At each iteration, the algorithm evaluates all possible combinations of maintenance actions for each state of the infrastructure system. The immediate reward (derived from maintenance and user-cost considerations, then transformed via the CRRA utility) is first transformed via the CRRA utility function, then combined with the discounted future value to determine the optimal action for each system state. For every potential next state resulting from an action, the algorithm calculates the utility-adjusted reward to represent the decision-maker's risk preference. Formally, for a given state s and action a, the value function $Q(s,a)$ is computed as:

$$Q(s,a) = \sum_{s'} P(s'|s,a).\left[U\big(R(s,a,s')\big) + \beta\,.\,V(s')\right] \tag{2}$$

Where $P(s' \mid s,a)$ represents the probability of transitioning from state s to s' after applying action a, $U(R(s,a,s'))$ is the CRRA utility applied to the immediate reward (reflecting maintenance and user costs), and β is the discount factor. The action a^* that maximizes the expected value is selected according to:

$$a^* = \arg\max_{a \in A} Q(s,a) \tag{3}$$

The value function for state s is updated as:

$$V_{new}(s) = Q(s,a^*) \tag{4}$$

This process continues iteratively, with the value function updated until convergence. The maximum change across all states is calculated as:

$$\delta = \max_{s \in S} |V_{new}(s) - V(s)| \tag{5}$$

When δ falls below a predetermined threshold θ, the algorithm terminates, indicating that an optimal policy has been found.

The expected value of each action is then computed as the sum of the utility-weighted immediate reward and the discounted future value of the subsequent state. This approach ensures that the optimization not only seeks to minimize costs but also aligns with the risk attitudes of infrastructure managers. By varying the risk aversion parameter γ, the framework yields policies spanning risk-averse, neutral, and seeking behaviours; in turn, this risk-adjusted reward structure drives markedly different maintenance mixes across the planning horizon.

3. Results and discussions

3.1 Case Study Description

To demonstrate the proposed risk-aware maintenance optimization framework, a case study was developed using a representative section of the Dutch transportation infrastructure. The selected network includes two major motorways—the A13 and A4—and two critical river crossings, the Schie River Bridge and the Vliet River Bridge. This corridor forms a key connection between Rotterdam and The Hague and is characterized by high traffic volumes and significant strategic importance. The infrastructure configuration mirrors a series-parallel system layout, wherein the A13 motorway and Schie River Bridge form one branch, and the A4 motorway and Vliet River Bridge form a parallel alternative branch (Figure 1).

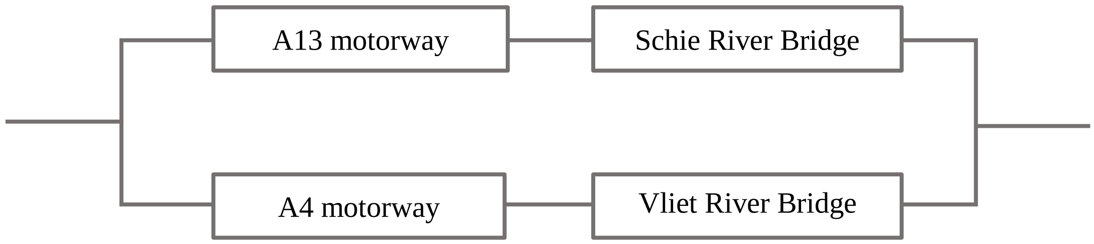

Fig. 1: Graph representation of the transportation network.

This arrangement ensures that traffic can continue to flow even when maintenance or deterioration affects one branch, aligning well with the study's focus on evaluating how risk attitudes influence maintenance decisions in redundant, multi-component systems. The layout of the system components is purposely chosen to allow clear visualization of the optimal behaviors, including preventive, corrective, opportunistic, and induced effects.

Each infrastructure component is represented by five discrete condition states, progressing from an as-new condition (state 0) to complete failure (state 4). The available maintenance actions comprise do-nothing, minor maintenance, and major maintenance, with each action defined by distinct cost implications, downtime requirements, and probabilistic state transition matrices. The deterioration and maintenance transition matrices for the network components are adopted from previous validated studies. For road segments, the matrices are sourced from Fathy et al. (2025), which provides deterioration and maintenance models calibrated for highway pavement systems [1]. For bridges, the matrices are taken from Faddoul et al. (2011) [12]. The adopted transition matrices are presented in Table1.

Table 1: Maintenance Action / Deterioration Transition Matrices.

Element	**Deterioration**	**Maintenance 1**	**Maintenance 2**
Roads	$\begin{bmatrix} 0.85 & 0.15 & 0 & 0 & 0 \\ 0 & 0.9 & 0.1 & 0 & 0 \\ 0 & 0 & 0.75 & 0.25 & 0 \\ 0 & 0 & 0 & 0.6 & 0.4 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0.85 & 0.15 & 0 & 0 & 0 \\ 0 & 0.9 & 0.1 & 0 & 0 \\ 0.2 & 0 & 0.6 & 0.2 & 0 \\ 0.05 & 0 & 0 & 0.57 & 0.38 \\ 0.02 & 0 & 0 & 0 & 0.98 \end{bmatrix}$	$\begin{bmatrix} 0.85 & 0.15 & 0 & 0 & 0 \\ 0 & 0.9 & 0.1 & 0 & 0 \\ 0.02 & 0 & 0.735 & 0.245 & 0 \\ 0.05 & 0 & 0 & 0.57 & 0.38 \\ 0.2 & 0 & 0 & 0 & 0.8 \end{bmatrix}$
Bridges	$\begin{bmatrix} 0.5 & 0.25 & 0.2 & 0.05 & 0 \\ 0 & 0.5 & 0.25 & 0.2 & 0.05 \\ 0 & 0 & 0.5 & 0.3 & 0.2 \\ 0 & 0 & 0 & 0.7 & 0.3 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0.9 & 0.1 & 0 & 0 & 0 \\ 0.15 & 0.7 & 0.15 & 0 & 0 \\ 0.1 & 0.3 & 0.5 & 0.1 & 0 \\ 0.4 & 0.3 & 0.2 & 0.1 & 0 \\ 0.2 & 0.4 & 0.2 & 0.1 & 0.1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0.7 & 0.3 & 0 & 0 & 0 \\ 0.4 & 0.4 & 0.2 & 0 & 0 \\ 0 & 0.2 & 0.3 & 0.4 & 0.1 \\ 0 & 0 & 0.3 & 0.3 & 0.4 \end{bmatrix}$

The maintenance action costs and downtimes are summarized in Table2.

Table 2 : Maintenance Actions Cost / Downtime.

Element	**Maintenance1 Cost (units)**	**Maintenance1 Downtime**	**Maintenance2 Cost (units)**	**Maintenance2 Downtime**
Roads	200	0.1	600	0.4
Bridges	300	0.5	800	0.8

The value iteration algorithm incorporating the CRRA utility function successfully converged to optimal policies for the series-parallel network across the tested risk preference scenarios. In this study, the coefficient of relative risk aversion γ was systematically varied with the following values: 0, 0.5, 1, 2, 3, and 4, thereby covering a spectrum from risk-neutral ($\gamma = 0$), through moderately risk-averse ($\gamma = 1$), to highly risk-averse ($\gamma \geq 2$) attitudes. The convergence of the value function was achieved within a reasonable number of iterations for all risk settings, confirming the computational efficiency and stability of the proposed reinforcement learning framework.

3.2 Total Cost Comparison Across Risk Aversion Parameter

The optimal policies derived under different risk preferences resulted in distinct total cost outcomes, reflecting how decision-makers' attitudes toward risk influence maintenance planning. Figure 2 presents the total expected discounted costs for various representative system states under different values of the risk aversion coefficient (γ = 0, 0.5, 1, 2, 3, and 4). As the risk aversion coefficient γ increases, the model consistently prioritizes preventive maintenance strategies, which initially incur higher short-term expenditures but substantially reduce long-term user costs by avoiding costly system failures and extensive downtimes.

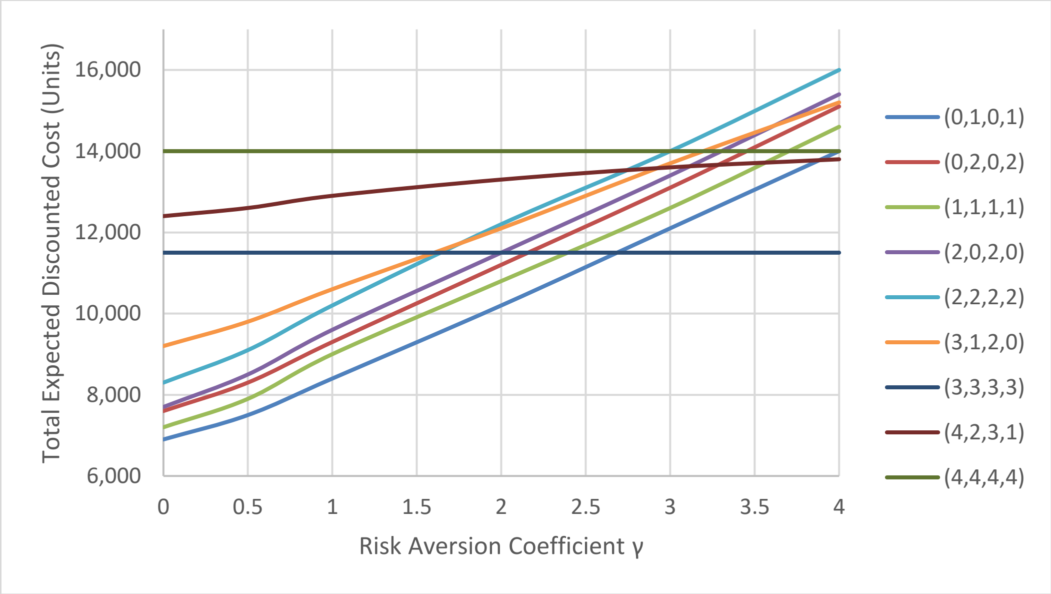

Fig. 2: Total expected discounted costs versus risk aversion coefficient γ for selected infrastructure system states.

The graphical results clearly illustrate that increased risk aversion significantly increases the overall expected costs, particularly for moderately deteriorated system states. For instance, the total discounted cost for a moderately deteriorated state such as (1,1,1,1) rises from approximately 7,200 units under a risk-neutral approach ($\gamma = 0$) to about 14,600 units under a highly risk-averse scenario ($\gamma = 4$), reflecting intensified proactive maintenance to prevent further deterioration. Conversely, for highly deteriorated states such as (4,4,4,4), total expected costs remain largely unaffected by risk attitude variations, since corrective maintenance becomes mandatory regardless of the adopted risk preference.

It is worth noting that, while much of the literature associates higher risk aversion with lower lifecycle costs—owing to the avoidance of catastrophic failures—the present study reveals the opposite trend. In this case, the emphasis on early preventive actions under high γ leads to higher total discounted costs. This outcome highlights the economic trade-off between uncertainty reduction and upfront expenditures, underscoring how the valuation of preventive interventions critically shapes the overall cost profile.

3.3 Analysis of Optimal Policy Behaviors

The optimal policies revealed clear behavioral patterns aligned with the decision-makers' risk attitudes. Beyond well-established preventive, corrective, and opportunistic actions, this analysis reveals the emergence of induced maintenance, a system-driven phenomenon that arises in series-parallel networks where interventions on one branch necessitate or trigger maintenance on alternative routes to ensure uninterrupted service delivery.

Induced maintenance is defined as maintenance performed on a component or branch (for example, Branch 2 or an alternative road) not solely due to its own condition, but as a strategic response to planned interventions or deteriorating conditions on another branch (e.g., Branch 1 or the main route). This approach leverages the redundancy inherent in parallel infrastructure to coordinate interventions and mitigate the risk of simultaneous failures or excessive downtime.

The influence of risk preferences on the distribution of maintenance actions is quantitatively presented in Figure 3. As the risk aversion parameter γ increases, the proportion of preventive maintenance rises markedly, while corrective actions decline. Notably, induced maintenance also becomes more prominent under risk-averse settings, reflecting a strategic shift toward coordinated, system-level interventions that proactively manage interdependencies and reduce network vulnerability.

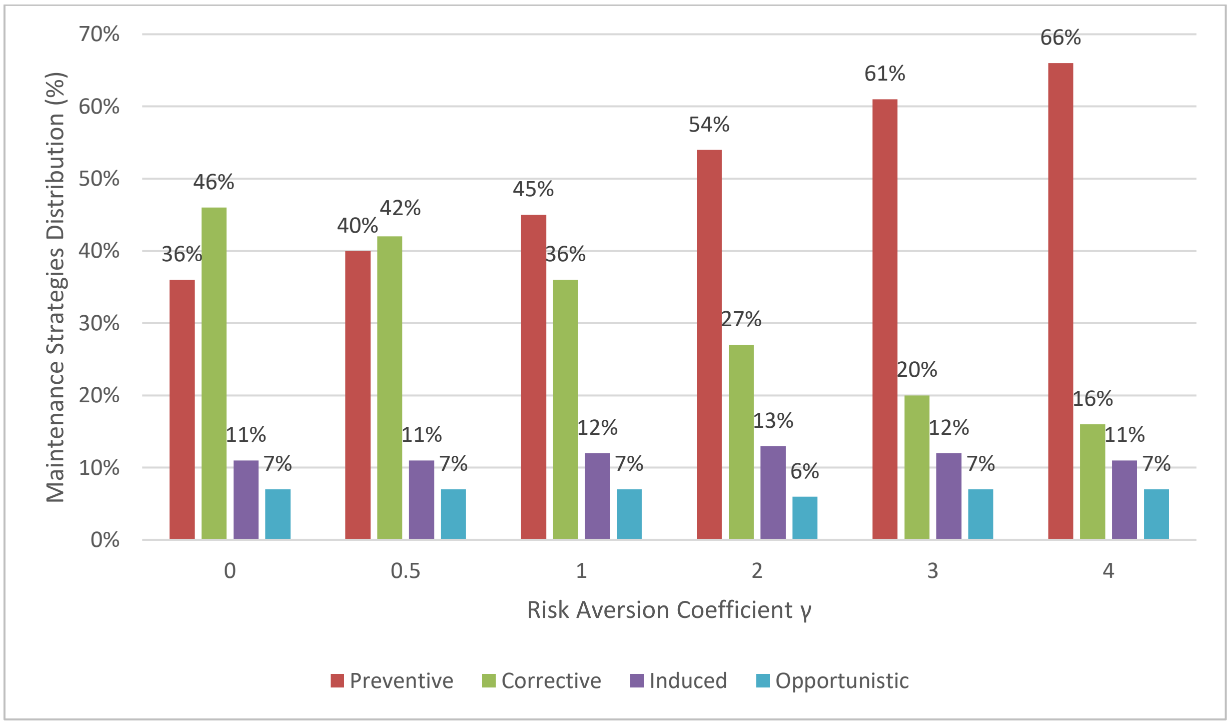

Fig. 3: Influence of risk aversion parameter (γ) on the distribution of maintenance strategies in the optimal policy.

These findings underscore the critical role of risk-informed strategies—not only in optimizing individual component reliability, but in orchestrating interventions across parallel branches to enhance overall system resilience and operational continuity.

4. Conclusion

This study introduces a risk-aware optimization framework for maintenance planning in transportation infrastructure systems, explicitly integrating decision-makers' risk attitudes into the strategy selection process. By extending a value iteration reinforcement learning algorithm with the Constant Relative Risk Aversion (CRRA) utility function, the framework enables the modeling of risk-averse behaviors within maintenance decision-making. The approach is demonstrated on a representative series-parallel network consisting of major motorways and bridges within the Dutch transportation infrastructure. Using validated deterioration and maintenance transition matrices, the model successfully generates optimal maintenance policies under varying risk preferences.

The results highlight how different risk attitudes significantly affect the selection and timing of maintenance strategies, balancing preventive, corrective, opportunistic, and induced maintenance in distinct ways. Risk-averse policies prioritize early preventive maintenance, resulting in higher preventive and lower corrective maintenance proportions as γ increases. Risk-neutral and risk-seeking strategies, in contrast, delay interventions and rely more on corrective maintenance. Interestingly, while much of the literature reports that higher risk aversion reduces lifecycle costs by avoiding catastrophic failures, this study shows that overall costs increase under high γ. This occurs because preventive maintenance is prioritized even in slightly degraded states, whereas for heavily degraded states corrective maintenance dominates regardless of risk attitude. The proportions of opportunistic and induced maintenance actions remain relatively stable across different risk attitudes, though the induced effect continues to play a relevant role, particularly in more complex system configurations.

The framework's ability to tailor maintenance policies based on organizational risk preferences strengthens its applicability for infrastructure managers, offering flexible and realistic decision-support tools. It aligns advanced optimization techniques with the practical complexities of infrastructure management, where diverse risk tolerances and resource constraints influence strategic choices. Future work should investigate uncertainty in deterioration and cost parameters to evaluate how robust the proposed risk-aware policies remain under variable real-world conditions.

References

[1] A. E. A. E. R. Fathy, M. Abdel-Hamid, and H. M. Abdelhaleem, “Developing a predictive maintenance model for the rehabilitation of roads’ assets,” *Int. J. Constr. Manag.*, vol. 25, no. 4, pp. 398–408, 2025.

[2] W. Zhang, D. Yang, and H. Wang, "Data-driven methods for predictive maintenance of industrial equipment: A survey", *IEEE Syst. J.*, vol. 13, no. 3, pp. 2213-2227, 2019.

[3] A. M. Smith, *Reliability-Centered Maintenance*. New York: McGraw-Hill, 1993, p. 45.

[4] C. R. Cassady, E. A. Pohl, and W. P. Murdock, "Selective maintenance modeling for industrial systems", *J. Qual. Maint. Eng.*, vol. 7, no. 2, pp. 104-117, 2001.

[5] R. Ahmad and S. Kamaruddin, "An overview of time-based and condition-based maintenance in industrial application", *Computers & Ind. Eng.,* vol. 63, no. 1, pp. 135-149, 2012.

[6] A. Grall, L. Dieulle, C. Bérenguer, and M. Roussignol, "Continuous-time predictive-maintenance scheduling for a deteriorating system", *IEEE Trans. Rel.*, vol. 51, no. 2, pp. 141-150, 2002.

[7] J. W. Pratt, “Risk aversion in the small and in the large,” in *Uncertainty in Economics*, Academic Press, pp. 59–79, 1978.

[8] M. M. Claramunt, M. Mármol, and X. Varea, “Facing a risk: To insure or not to insure—An analysis with the constant relative risk aversion utility function,” *Mathematics*, vol. 11, no. 5, p. 1070, 2023.

[9] K. A. Kobbacy and D. N. Murthy, *Complex System Maintenance Handbook*. New York, NY, USA: Springer, 2008.

[10] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Hoboken, NJ: John Wiley & Sons, 2014.

[11] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. Cambridge, MA, USA: MIT Press, 1998.

[12] R. Faddoul, W. Raphael, and A. Chateauneuf, “A generalised partially observable Markov decision process updated by decision trees for maintenance optimisation,” *Struct. Infrastruct. Eng.*, vol. 7, no. 10, pp. 783–796, 2011.