Modelling stochastic degradation and maintenance effects for the road network of Amsterdam: a multi-attribute data-driven approach

Sterrenberg, A.G.E.; Andriotis, C.; Stoter, J.E.

DOI
10.23967/icossar.2025.099

Publication date
2025

Document Version
Final published version

Citation (APA)
Sterrenberg, A. G. E., Andriotis, C., & Stoter, J. E. (2025). *Modelling stochastic degradation and maintenance effects for the road network of Amsterdam: a multi-attribute data-driven approach*. Paper presented at 14th International Conference on Structural Safety and Reliability, Los Angeles, United States. https://doi.org/10.23967/icossar.2025.099

14th International Conference on Structural Safety and Reliability - ICOSSAR'25
1-6 June, 2025 - Los Angeles California, USA
R. Ghanem, E. Johnson, S. Masri and A. Olivier (Eds)

MODELLING STOCHASTIC DEGRADATION AND MAINTENANCE EFFECTS FOR THE ROAD NETWORK OF AMSTERDAM: A MULTI-ATTRIBUTE DATA-DRIVEN APPROACH

AMY G.E. STERRENBERG, CHARALAMPOS P. ANDRIOTIS, JANTIEN E. STOTER

Faculty of Architecture and the Built Environment, Delft University of Technology
2628 BL Delft, The Netherlands
A.G.E.Sterrenberg@tudelft.nl, C.Andriotis@tudelft.nl, J.E.Stoter@tudelft.nl

Key words: Stochastic degradation modelling; Maintenance impact estimation; Life-cycle extension; Predictive maintenance; Data-driven infrastructure management; Markov models

Abstract. Data-driven prediction of infrastructure aging is challenging due to the complex stochastic nature of degradation effects and the ill-documented historical records. Degradation modeling is, however, crucial for predictive maintenance that is key for infrastructure integrity. This study presents a multi-attribute, data-driven approach for modelling stochastic degradation and maintenance effects of roads, mining an extensive database of geo-located historical inspection and maintenance records from the municipality of Amsterdam. Inspection data track pavement conditions at irregular intervals across ten discrete states per road segment, following the Dutch CROW 146 protocol. Damage severity and extent for eight damage modes is captured, i.e., for transverse unevenness, irregularities, ravelling, edge damage, crack formation, joint filling, joint width, and settling. The maintenance dataset includes >25k minor interventions across 17k road segments, indicating repair requirements, and 200+ major maintenance projects, covering 21k segments where interventions are planned, all without verifying completion. This complicates accurate modelling of natural degradation as it is confounded by maintenance effects. To address the issue of irregular inspections, degradation is first modelled as a continuous-time Markov chain. Thereby, transition rates are estimated, which are then converted to discrete-time Markov chain transition probability matrices to eventually support regular maintenance planning. We further learn the effects of minor and major maintenance activities, as defined and recorded in the database. Based on the estimated degradation transitions, pre-maintenance and post-maintenance state distributions are estimated. Instantaneous maintenance transition matrices are computed by minimizing the cross-entropy between the pre-maintenance state after the intervention and the post-maintenance state. The model allows for a multi-attribute approach, segmenting roads based on construction material (e.g., asphalt, tiled pavement) and traffic loads (e.g., residential, commercial/pedestrian). The approach is exemplified for tiled pavements for a section of the road network of Amsterdam, where the effects of minor and major maintenance are ablated for long-term predictions. Although applied to Amsterdam, this method is relevant to any infrastructure system with discrete state datasets, providing a foundation for data-driven decision-making in infrastructure management.

1 INTRODUCTION

Urban transportation infrastructure is critical for mobility, economic activity, and the serviceability of cities. The maintenance of these systems is essential to preserve operational reliability and safety. However, infrastructure management involves complex, multi-objective decision-making under resource constraints. Authorities do not only prioritise interventions based on urgency and cost-effectiveness, but also consider broader system-level impacts such as network disruptions, emissions, accessibility, and safety risks. To control these impacts as demands on infrastructure increase, efficient resource allocation becomes more and more crucial.

Reactive, unplanned or corrective maintenance, where damage is addressed as it is observed, remain common practice by operating agencies. While these approaches minimise short-term expenditures, they are generally associated with higher overall maintenance costs, increased risk of asset failure, and reduced system reliability in the long run [1, 2, 3, 4]. In contrast, predictive maintenance (PdM) offers a more proactive alternative, using condition monitoring and data analysis to anticipate infrastructure needs and optimise intervention timings [1, 3, 5].

Accurate degradation modelling facilitates the development of effective PdM decision-making systems. While data quality and availability have been limiting factors in the past, data-driven degradation models have become more and more prevalent and have been extensively explored for various applications, using methods such as statistical inference, probabilistic graphs, including Markov models, and machine learning [6]. Similar models have been developed for pavements (e.g., [7, 8, 9, 10, 11, 12, 13, 14, 15]). These models are typically calibrated using indicators of pavement condition, such as the pavement condition index (PCI), international roughness index (IRI), pavement serviceability index (PSI), or pavement condition rating (PCR). Such continuous indices allow modeling through continuous-state stochastic models such as gamma processes [16]. Additionally, most studies focus on asphalt and concrete pavements, for which these indicators are available, limiting their applicability to urban road networks characterised by heterogeneous attributes, such as diverse pavement materials, traffic loads, damage modes and other functional classifications [17]. Moreover, models explicitly quantifying the effect of maintenance activities on road condition are currently limited. Although some studies have explored post-maintenance condition changes (e.g., [18, 19]), there is limited integration between degradation and maintenance modelling in a unified probabilistic framework. This omission constrains the accuracy of long-term maintenance planning and limits the potential for rigorous predictive optimisation methods, such as Markov decision processes and reinforcement learning, to be applied in practice [20, 16, 21, 22].

This study presents a data-driven, multi-attribute framework for the modelling of both road degradation and maintenance effects, designed to support predictive maintenance in heterogeneous road networks present in urban environments. The model is developed using inspection, maintenance and road data from the municipality of Amsterdam, the Netherlands. The inspection data includes information on the severity and extent of damage across up to eight predefined damage modes, with up to six damage modes depending on pavement construction type. The methodology allows for condition transitions to be modelled as probabilistic state changes conditioned on segment-specific attributes, including construction type and traffic pattern categories. These transitions are initially formulated as a continuous-time Markov chain (CTMC) to address the irregular time intervals present in the observational data and are subsequently converted into a discrete-time Markov chain (DTMC) to support integration with regular, time-stepped maintenance planning frameworks. The resulting transitions form a high-dimensional tensor, which enables detailed, scenario-based simulation of network condition under varying intervention policies. As a result, the proposed framework supports future integration with decision-support tools for high-fidelity predictive maintenance planning at scale, offering an interpretable and empirically grounded basis for modelling deterioration of diverse road networks. Results regarding future prediction in an area of Amsterdam are discussed, ablating minor and major maintenance scenarios.

2 METHODS

2.1 Data

Historical inspection and maintenance data provided by the municipality of Amsterdam are used. The inspection dataset contains manually collected road condition data in accordance with the CROW 146b protocol, which is the national standard for infrastructure inspections in the Netherlands. Inspections are conducted per road segment, defined as a continuous section of road between intersections with a maximum length of 100 metres. The dataset is primarily used by the municipality for maintenance planning, administrative reporting, and compliance with regulatory obligations.

Inspection records have been collected since approximately 2015, with road segments inspected at irregular time intervals, mostly every two to three years. Each inspection record includes a condition label that reflects the most severe observed damage within the segment. These labels—G, $L1$–$L3$, $M1$–$M3$, and $E1$–$E3$—represent a combination of severity (G = Good, L = Light, M = Moderate, E = Serious) and extent of the affected area (1 = small, 2 = average, 3 = large).

Eight damage modes are used in the classification: transverse unevenness; irregularities; ravelling; edge damage; crack formation; joint filling; joint width; and settling. Up to six unique damage modes are relevant per pavement category: asphalt segments exhibit transverse unevenness, irregularities, ravelling, edge damage, cracking, and settling; tiled pavements show transverse unevenness, irregularities, joint width, and settling; and concrete surfaces are assessed for irregularities, cracking, joint filling, and settling. However, no or very limited degradation is represented in inspection records for settling and joint filling. Although the CROW protocol defines boundary conditions for each label and damage mode, expert opinion suggests a degree of human error in the dataset.

Additional features are available for each road segment, including geographical information (i.e., neighbourhood, region and geometry), construction type (e.g., asphalt, concrete, element pavement, semi-paved, synthetic, unpaved), specific surface material, year of construction or last major reconstruction, date of last conservation treatment, surface area, and traffic pattern classification (e.g., heavy, residential, cycle path). The completeness of this data varies across records.

In addition to inspection records, two separate datasets of maintenance records are used: one for minor interventions and one for major interventions. The minor maintenance dataset contains 25,044 intervention points identified by road inspectors, which we can link to 17,813 unique road segments. Each record includes coordinates, date of data collection, pavement type, damage category, damage size, and corresponding unit of measurement. The major maintenance dataset includes 228 planned interventions covering larger areas, represented spatially as polygons or multi-polygons. These polygons typically include multiple road segments. Each record documents the date of data collection, project name, category, budget, asset group, source, project status, and planned phases of execution: preparation start, intervention start, and completion date, each recorded by year. The total area scheduled for maintenance is also included.

2.2 Modelling Approach

In this study, both the degradation and maintenance effects are modelled through probabilistic Markovian transition matrices, derived from empirical data. This section first outlines the fitting of degradation probability matrices, followed by a description of how minor and major maintenance effects are incorporated.

Degradation Modelling

Degradation of road infrastructure is modelled as a stochastic process using a continuous-time Markov chain (CTMC). This model represents transitions between a finite set of condition states in continuous time, which evolve over time depending on discrete actions. The probability of transitioning to a future state adheres to the Markov property, meaning that the probability of moving to a future state depends solely on the current state and not on the sequence of preceding states. The CTMC model is based on several key assumptions. First, transitions between states are assumed to be independent across road segments; that is, the degradation of one segment does not influence others. Second, the model assumes stationarity of transition rates, implying that these remain constant over time.

A Markov process is defined as $\{X(t),\, t \geq 0\}$ on a finite state space $S = \{1, 2, \ldots, n\}$. In this work, the state space has $n = 10$ states, corresponding to the inspection labels G–$E3$, described in Section 2.1. The transition between states is characterized by a transition probability matrix, where the probability of transitioning from state $i \in S$ to state $j \in S$ in exactly time Δt from an initial time t is denoted as [23]:

$$P_{ij}(t) = P(X(t + \Delta t) = j | X(\Delta t) = i) \tag{1}$$

This defines the transition probability matrix $P(t)$ at time t:

$$P(t) = \begin{bmatrix} P_{11}(t) & P_{12}(t) & \cdots & P_{1n}(t) \\ P_{21}(t) & P_{22}(t) & \cdots & P_{2n}(t) \\ \vdots & \vdots & \ddots & \vdots \\ P_{n1}(t) & P_{n2}(t) & \cdots & P_{nn}(t) \end{bmatrix} \tag{2}$$

Each row of $P(t)$ sums up to 1, and all entries satisfy $0 \leq P_{ij}(t) \leq 1$. The transition probability matrix for timestep t is used in discrete-time Markov processes, where $P_{ij}(t)$ can be estimated empirically by the observed transition frequencies [24, 25, 26]:

$$P_{ij}(t) = \frac{n_{ij}}{n_i} \tag{3}$$

where n_{ij} is the number of transitions from state i to state j over the observation period, and n_i is the total number of transitions from state i to any other state in the state space. However, given that the available inspection data in this study contains variable time intervals, a discrete-time approach is not suitable. Therefore, we employ CTMCs to model the degradation process. In this case, transition probability matrices are derived from transition rate matrices. Assuming exponentially distributed sojourn and transition times, the transition rate matrix Q takes the form:

$$Q = \begin{bmatrix} -\lambda_1 & \lambda_{12} & \cdots & \lambda_{1n} \\ \lambda_{21} & -\lambda_2 & \cdots & \lambda_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_{n1} & \lambda_{n2} & \cdots & -\lambda_n \end{bmatrix} \tag{4}$$

Following Eq 1, each off-diagonal entry λ_{ij} for $i \neq j$ is defined as the rate of transition from state i to state j [27, 28]:

$$\lambda_{ij} = \lim_{\Delta t \to 0^+} \frac{P(X(\Delta t) = j \mid X(0) = i)}{\Delta t} \tag{5}$$

The diagonal elements λ_i in the transition rate matrix Q reflect the total rate at which the segment leaves state i, and are defined such that each row of Q sums to zero:

$$\lambda_i = \sum_{j \neq i} \lambda_{ij} \tag{6}$$

Under the exponential distribution assumptions, the transition rates λ_{ij} are estimated from the observed transitions in the inspection dataset. The estimate for each λ_{ij} is given by [27, 29]:

$$\lambda_{ij} = \frac{n_{ij}}{\sum_{k=1}^{n_i} T_i^{(k)}} \tag{7}$$

where: $T_i^{(k)}$ is the observed sojourn time in state i before transitioning, for the k-th transition. To obtain the corresponding transition probability matrix $P(t)$ for a desired time horizon t, we compute the matrix exponential [23]:

$$P(t) = e^{Qt} \tag{8}$$

This equation allows us to derive discrete-time transition probability matrices, which can subsequently be used in advanced decision-making algorithmic frameworks for predictive maintenance (Figure 1), such as a Markov decision process or within (deep) reinforcement learning.

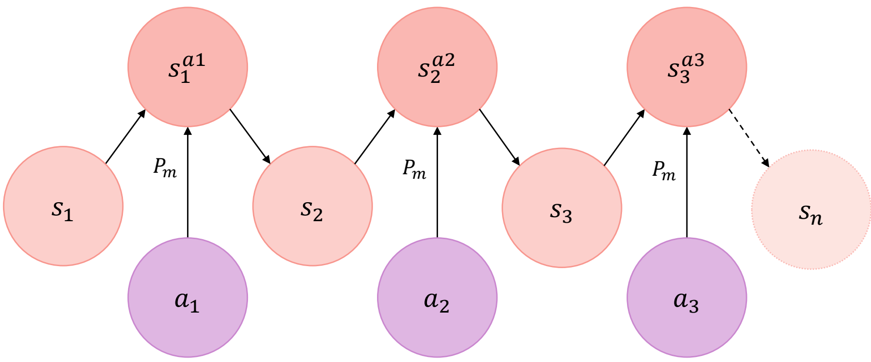

Figure 1: Probabilistic graph of the DTMC where the state space $S = \{s_1, s_2, \ldots, s_n\}$. Arrows indicate degradation transitions and P_m is the maintenance transition model.

Maintenance Effect Modelling

Maintenance effects are modelled separately for minor and major maintenance interventions. Transitions for which a maintenance action is recorded are considered in this process. Condition improvements without maintenance logs are excluded, as we lack sufficient data to interpret them reliably. As maintenance records include planned interventions, without confirmation of execution, the dataset may contain cases of failed, postponed, or unperformed maintenance. Additionally, some maintenance-logged transitions result in deterioration of condition. These cases can be interpreted as either: (1) failed maintenance, (2) maintenance unrelated to the observed damage mode, or (3) rapid post-maintenance degradation. Each inspection record includes the condition of the segment for multiple damage modes, dependent on road construction category. We assume all damage modes are subject to the maintenance intervention, regardless of the resulting change in condition.

The effect of a maintenance action taken at time t_m is modelled as an instantaneous transition probability matrix P_m, representing the maintenance effect between the time just before the maintenance intervention, t_m^-, and the time immediately after, t_m^+. Using the estimated degradation transition rate matrix Q_d from Eq.6 and the state probability distributions at the start, $p(t_{\text{start}})$, and end, $p(t_{\text{end}})$, of the observed transition, the predicted state vector just before maintenance is computed as:

$$p(t_m^-) = p(t_{\text{start}}) \cdot e^{Q_d(t_m - t_{\text{start}})} \tag{9}$$

The predicted state just after maintenance is computed through a backward transition as:

$$p(t_m^+) = p(t_{\text{end}}) \cdot e^{-Q_d(t_{\text{end}} - t_m)} \tag{10}$$

For major maintenance, the exact intervention date t_m is unknown. We assume it occurs at the midpoint of the time interval between inspections. Next, P_m is estimated for k observed maintenance events so that:

$$\begin{bmatrix} p(t_m^-)^{(1)} \\ p(t_m^-)^{(2)} \\ \vdots \\ p(t_m^-)^{(k)} \end{bmatrix} \cdot P_m = \begin{bmatrix} p(t_m^+)^{(1)} \\ p(t_m^+)^{(2)} \\ \vdots \\ p(t_m^+)^{(k)} \end{bmatrix} \tag{11}$$

We assume that P_m in Eq.11 has a known shape but unknown parameters. The shape is given by a linear combination between a perfect repair and a no-repair action. These transition matrices are denoted as P_m^1 and P_m^2, respectively. P_m^1 represents a transition probability matrix where $p(s' = G \mid s) = 1$ for all $s \in S$. This indicates that, regardless of the current state s, the system deterministically transitions to state G, effectively resetting the system to the best condition. P_m^2 is the identity matrix I over the state space S, implying that no transition occurs and the system remains in its current state. As such, the instantaneous transition for each damage mode assumes the following form:

$$P_m = w_1 \cdot [p(s' = G \mid s)] + w_2 \cdot [I], \quad w_1 + w_2 = 1 \tag{12}$$

Subsequently, we optimize the weights w_1 and w_2 based on minimizing the cross-entropy loss between predicted and observed post-maintenance state distributions across all N observed transitions:

$$L = \frac{1}{N} \sum_{j=1}^{N} \left(- \sum_{i=1}^{n} p_i^{(j)}(t_m^+) \log \hat{p}_i^{(j)}(t_m^+) \right) \tag{13}$$

where $p_i^{(j)}(t_m^+)$ is the probability of sample j being in state i at time t_m^+ based on Eq 10 and $\hat{p}_i^{(j)}(t_m^+)$ is the predicted probability of sample j being in state i at time t_m^+ based on Eq 11.

3 RESULTS

3.1 Data Selection

The dataset comprises comprehensive inspection records, corresponding to eight unique damage modes with up to six distinct damage types per road segment, in addition to feature data concerning road construction materials and traffic pattern classifications. This level of detail facilitates the specification of degradation and maintenance effects across various combinations of material and traffic pattern attributes. However, the dataset does not provide sufficient coverage to enable robust modelling of all such combinations.

With regard to construction materials, tiled pavements and asphalt pavements constitute the majority of the network. Specifically, 80.1% of roads in Amsterdam are classified as tiled pavements, and 15.9% as asphalt pavements. All remaining material categories individually represent less than 3% of the road network. While further disaggregation of the transition model (e.g., by traffic pattern classification) is possible, for the sake of brevity, aggregated results for tiled pavements across all traffic pattern classifications and for damage modes ’transverse unevenness’, ’irregularities’ and ’joint width’ are presented here.

3.2 Degradation Modelling

Degradation processes are modelled by estimating transition probability matrices derived from observed condition transitions. Figure 2 presents the 8-year transition probability matrices for the three dominant damage modes observed in tiled pavements within the Amsterdam road network: transverse unevenness, irregularities, and joint width. Each matrix displays the starting states along the vertical axis and the corresponding ending states along the horizontal axis. States are coded as follows: G (good/no damage), $L1 - L3$ (minor damage), $M1 - M3$ (moderate damage), and $E1 - E3$ (serious damage), with numerical indices denoting the spatial extent of damage (1 = small, 2 = average, 3 = large).

Table 1 summarises the total number of recorded transitions for each damage mode, as well as the proportion of observations in which no change in condition occurred ($G \rightarrow G$). High proportions

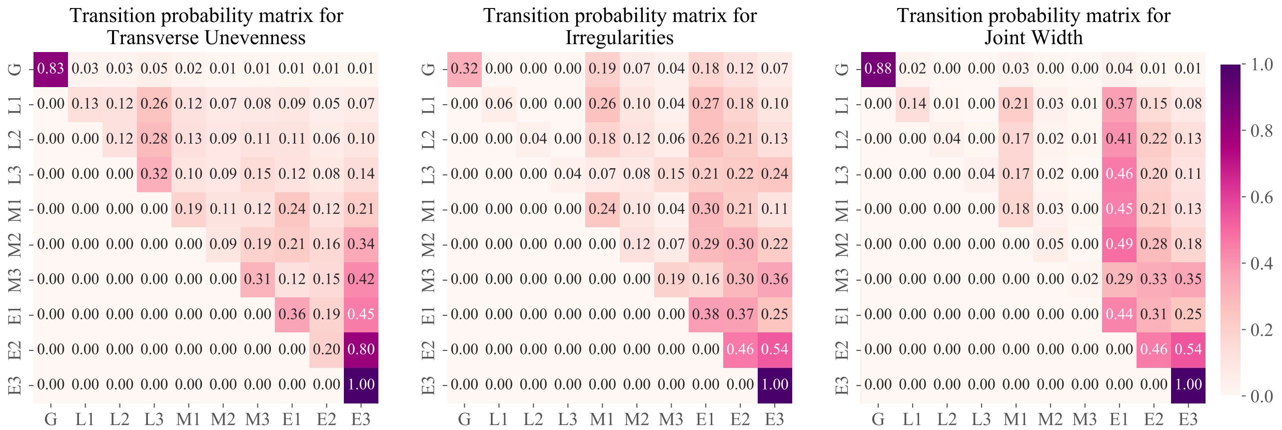

Figure 2: Eight-year transition probability matrices for the three relevant damage modes in tiled pavements across the Amsterdam road network: transverse unevenness, irregularities, and joint width. Condition states range from G (good/no damage) to E3 (severe damage affecting a large area), with severity increasing from top to bottom and left to right.

Table 1: Logged transitions and proportion of 'G' → 'G' transitions by damage mode

Damage Mode	Number of logged transitions	Number of 'G' → 'G' transitions
Transverse unevenness	73,249	67,087 (91.6%)
Irregularities	60,086	22,878 (38.1%)
Joint width	78,577	74,220 (94.5%)

of such stable states, particularly in joint width (94.5%) and transverse unevenness (91.6%), suggest relatively slow deterioration processes for those damage types. In contrast, the proportion of unchanged states is considerably lower for irregularities (38.1%), indicating more dynamic or variable degradation behaviour.

Analysis of the dataset shows that some condition states, particularly those representing more severe or extensive damage, are sparsely populated. As such, the reliability of transition probability estimates from these states, specifically $M3$, $E2$, and $E3$, is limited. Notably, for irregularities, the starting states $L2$ and $L3$ are also especially underrepresented. To isolate degradation processes, transitions reflecting stable or worsening conditions are included in the estimation. The most severe state ($E3$) is treated as an absorbing or terminal state in this model.

The values along the diagonal of each matrix indicate the probability that, over an eight-year period, segments remain in their current condition state. These values are especially high for $G \rightarrow G$ transitions for the damage modes 'transverse unevenness' and 'joint width'. Diagonal values are also higher for states indicating high damage severity ($E1$–$E3$), with an exception for $E2$ in transverse unevenness. For transverse unevenness, however, state transitions $L3 \rightarrow L3$ and $M3 \rightarrow M3$ are also relatively high. Non-zero values in the upper triangle suggest a gradual risk of deterioration. In particular, transitions from any state except G to $M1$ and $E1$ appear relatively likely for 'irregularities', and 'joint width' damage types. This trend is especially pronounced for joint width, suggesting that once deterioration initiates, it escalates more rapidly in this damage mode. Zero entries in the upper triangle of the matrix imply the absence of observed transitions to these states. However, although certain states are rounded down to zero in the matrices presented in Figure 2, the associated transition rates indicate existence of small probabilities. Finally, for transverse unevenness and irregularities, some transitions that indicate an increase in damage severity, while damage extent decreases or vice versa, such as $M3 \rightarrow E1$, are found to have slightly lower values compared to other transitions in the same row.

3.3 Maintenance Effect Modelling

Maintenance effects were modelled through the estimation of instantaneous transition probability matrices, derived from observed condition transitions that could be directly associated with either minor or major maintenance interventions. The weights of these matrices, estimated using observed transitions combined with a set of prototype matrices based on Eq. 12, are represented in Table 2. These describe the prototype matrices with the lowest cross-entropy relative to the empirical transition data, as outlined in the methodology section.

As dictated in Section 12, w_1 indicates how often maintenance interventions result in perfect repair (transition to state G) and w_2 indicates no observable change in condition. Results are shown in Table 3. Major maintenance consistently yields higher probabilities of perfect repair compared to minor maintenance. When comparing across damage modes, joint width exhibits the highest likelihood of

Table 2: Learned weights for maintenance impacts by damage mode and maintenance type.

Weight	Transverse Unevenness		Irregularities		Joint Width	
	Minor	Major	Minor	Major	Minor	Major
w_1	0.597	0.701	0.284	0.442	0.585	0.766
w_2	0.403	0.299	0.716	0.558	0.415	0.234

perfect repair, followed closely by transverse unevenness, whereas irregularities show substantially lower perfect repair rates. Additionally, for damage modes 'transverse unevenness' and 'joint width', the probability for perfect repair is always higher than the probability for no change in condition, indicating that maintenance interventions more commonly result in a perfect repair than no repair. For irregularities, maintenance actions have a smaller effect.

3.4 Application

Figure 3 illustrates a section of the Amsterdam road infrastructure network, to which the previously developed degradation and maintenance effect models have been applied. These comparative visualisations highlight the varying impact of maintenance strategies on the long-term condition of the network. The initial state of the network is shown in Figure 3a. Figure 3b illustrates the condition development over a eight-year period under natural degradation, without any intervention, indicating the condition state as the highest probability in the distribution over states for transverse unevenness. While some segments remain in their initial condition category, others exhibit progressive deterioration, particularly those that began in more severe damage states. Segments initially in perfect or

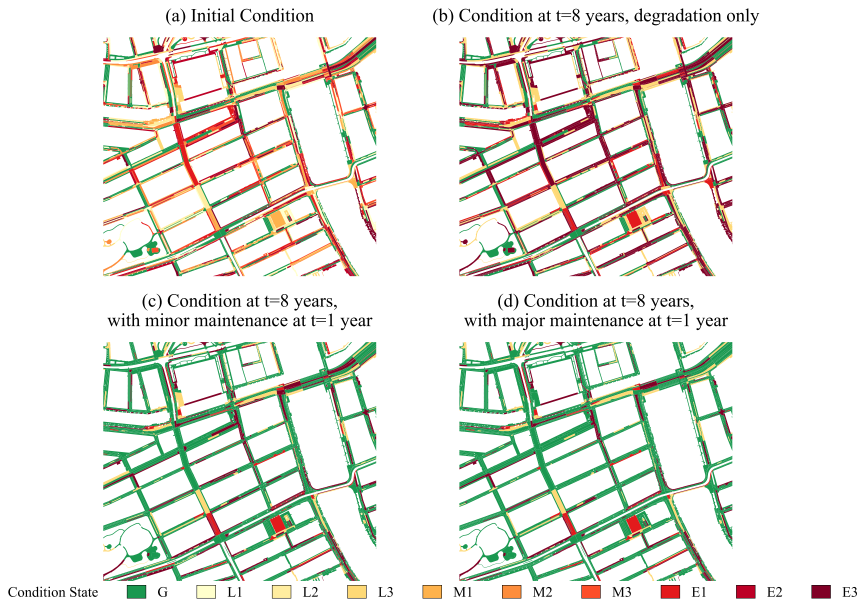

Figure 3: Visualisation of the condition evolution of a section of Amsterdam's road network under different maintenance scenarios, assuming tiled pavements for all segments: (a) Initial condition state; (b) Projected condition after eight years of natural degradation without intervention; (c) Condition after eight years with minor maintenance applied uniformly after one year; (d) Condition after eight years with major maintenance applied uniformly after one year. The comparison illustrates the differential impact of maintenance strategies on long-term infrastructure condition.

lightly deteriorated states generally show slower rates of degradation. Figure 3c shows the same scenario with minor maintenance applied uniformly to all road segments after year one, while Figure 3d reflects the scenario in which major maintenance is implemented at the same time point. In both cases, a substantial number of segments are restored to the optimal condition state; however, a subset of road segments shows no improvement following intervention. This proportion is notably higher under the minor maintenance scenario than under major maintenance, indicating the limited effectiveness of minor interventions in reversing deterioration in certain cases.

4 DISCUSSION

The data-driven model developed in this work enables degradation and maintenance effects transition modelling, which can serve as a foundation for predictive maintenance systems. The model captures the stochastic nature of deterioration and is grounded in condition data and classification schemes that align with those used by human maintenance planners. This makes it not only analytically robust but also potentially interpretable and actionable in operational contexts.

At the same time, several limitations exist with respect to the statistical analysis. First, the dataset was not originally collected for predictive modelling. It only records the most prominent manifestation of each damage mode and lacks precision regarding the exact timing and nature of maintenance interventions, as discussed in Section 2.1. This restricts our ability to draw strong conclusions about causal links between specific maintenance actions and changes in condition for particular damage types. As a result, observed condition improvements are associated with general maintenance categories (i.e., minor or major), as it is difficult to associate them with targeted treatments of, for example, transverse unevenness or joint width,

Another key constraint is data sparsity, particularly in higher-severity condition states. Some transitions, especially from rarely observed states, such as $E2$ or $E3$, are based on a few records, limiting the statistical confidence of the corresponding transition probabilities. This restricts the model's applicability in representing deterioration and maintenance effects across the full range of condition states.

Expanding the dataset to better cover underrepresented road types or damage states would improve the robustness and predictive power of the model, as well as increase its generalisability across the network. Despite these limitations, the modelling approach presented in this work demonstrates the potential of inspection records to support data-informed maintenance strategies. With improved data infrastructure, such models can significantly contribute to planning and prioritising maintenance interventions in urban road networks.

5 CONCLUSION

A transition model for infrastructure degradation and maintenance effects is presented. The model consists of transition probability matrices for a selection of road construction materials and traffic pattern labels and is aimed at applications in predictive maintenance models. Historical inspection and maintenance records from the municipality of Amsterdam were used, comprising data collected per road segment in accordance with the CROW 146b protocol. The inspection records include condition labels based on eight damage categories, together with additional features, such as road material, construction year, and traffic pattern category. Records of both minor and major maintenance activities are available at the same spatial resolution as the inspection data, enabling the modelling of condition state transitions under natural degradation, minor maintenance, and major maintenance. Degradation was modelled as a continuous-time Markov chain. This approach accounted for the irregular inspection intervals by estimating stationary transition rates from observations, assuming exponentially distributed transition times. Accordingly, transition probability matrices were derived through the exponential of the rate matrix. To address data uncertainties, including unexecuted or incorrectly logged maintenance, transition matrices were introduced, as a linear combination of perfect repair and no (or unsuccessful) repair. The maintenance effect matrix was estimated by comparing predicted and observed post-maintenance states, selecting the matrix weights by the lowest cross-entropy loss. Results were presented for tiled pavements and three damage modes: transverse unevenness; irregularities; and joint width. The degradation analysis showed that for the lowest and highest condition states, higher transition probabilities are generally found along the matrix diagonal. This suggests limited observed change in these states. Non-zero upper-triangular values, especially for transitions not originating from state G, indicate a comparatively faster deterioration once initiated. However, data sparsity in more severe condition states limits the reliability of certain estimates. Results also show that both minor and major maintenance commonly lead to perfect repair for 'transverse unevenness' and 'joint width' damage modes. For irregularities, maintenance actions were found more likely to not affect road condition. For all damage modes, major maintenance consistently delivers higher rates of full restoration. Among the three damage types, joint width was most responsive to maintenance, followed by transverse unevenness. Irregularities were less responsive. Overall, the results highlight differences in degradation pace and maintenance effectiveness across damage types and emphasise the need to account for these distinctions in infrastructure planning and maintenance strategies.

References

[1] Laura Swanson. "Linking maintenance strategies to performance". In: *Int. J. Prod. Econ.* 70.3 (2001), pp. 237–244.

[2] Keith R. Mobley. *An Introduction to Predictive Maintenance*. 2nd. Elsevier, 2002.

[3] P. Coandă, M. Avram, and V. Constantin. "A state of the art of predictive maintenance techniques". In: *IOP Conf. Ser. Mater. Sci. Eng.* Vol. 997. 1. 2020.

[4] Hikmet Erbiyik. "Definition of Maintenance and Maintenance Types with Due Care on Preventive Maintenance". In: *Maint. Manag. - Curr. Challenges, New Dev. Futur. Dir.* (2023).

[5] Mohammad M. Hamasha, Ala H. Bani-Irshid, Sahar Al Mashaqbeh, Ghada Shwaheen, Laith Al Qadri, Mohammad Shbool, Dania Muathen, Mussab Ababneh, Shahed Harfoush, Qais Albedoor, and Adnan Al-Bashir. "Strategical selection of maintenance type under different conditions". In: *Sci. Rep.* 13.1 (2023).

[6] Anna Jarosz, Marta Zagorowska, and Jerzy Baranowski. "Recent advances in data-driven methods for degradation modelling across applications". In: (2025).

[7] Rabi G. Mishalani and Samer M. Madanat. "Computation of Infrastructure Transition Probabilities Using Stochastic Duration Models". In: *J. Infrastruct. Syst.* 8.4 (2002), pp. 139–148.

[8] K. P. George, A. S. Rajagopal, and L. K. Lim. "Models for predicting pavement deterioration". In: *Transp. Res. Rec.* 1215 (1989), pp. 1–7.

[9] Turki I. Al-Suleiman and Adnan M.S. Shiyab. "Prediction of Pavement Remaining Service Life Using Roughness Data - Case Study in Dubai". In: *Int. J. Pavement Eng.* 4.2 (2003), pp. 121–129.

[10] Nader Abdelaziz, Ragaa T. Abd El-Hakim, Sherif M. El-Badawy, and Hafez A. Afify. "International Roughness Index prediction model for flexible pavements". In: *Int. J. Pavement Eng.* 21.1 (2020), pp. 88–99.

[11] Saeid Alimoradi, Amir Golroo, and Seyed Mohammad Asgharzadeh. "Development of pavement roughness master curves using Markov Chain". In: *Int. J. Pavement Eng.* 23.2 (2022), pp. 453–463.

[12] S. Madeh Piryonesi and Tamer E. El-Diraby. “Data Analytics in Asset Management: Cost-Effective Prediction of the Pavement Condition Index”. In: *J. Infrastruct. Syst.* 26.1 (2020).
[13] Ningyuan Li, Wei Chau Xie, and Ralph Haas. “Reliability-based processing of markov chains for modeling pavement network deterioration”. In: *Transp. Res. Rec.* 1524 (1996), pp. 203–213.
[14] Kiyoshi Kobayashi, Kiyoyuki Kaito, and Nam Lethanh. “A Bayesian Estimation Method to Improve Deterioration Prediction for Infrastructure System with Markov Chain Model”. In: *Int. J. Archit. Eng. Constr.* 1.1 (2012), pp. 1–13.
[15] Jhenyffer Lorrany Matias de Oliveira, Gary Davis, Alireza Khani, and Mihai Marasteanu. “Heterogeneous Markov Chain Model to Predict Pavement Performance and Deterioration”. In: *Transp. Res. Rec.* 2676.9 (2022), pp. 568–581.
[16] M. Saifullah, K. G. Papakonstantinou, C. P. Andriotis, and S. M. Stoffels. “Multi-agent deep reinforcement learning with centralized training and decentralized execution for transportation infrastructure management”. In: i (2024), pp. 1–42.
[17] Amir Shtayat, Sara Moridpour, Berthold Best, and Shahriar Rumi. “An Overview of Pavement Degradation Prediction Models”. In: *J. Adv. Transp.* 2022 (2022).
[18] Putu Mandiartha, Colin F. Duffield, Russell G. Thompson, and Marcus R. Wigan. “Measuring pavement maintenance effectiveness using Markov Chains analysis”. In: *Struct. Infrastruct. Eng.* 13.7 (2017), pp. 844–854.
[19] Yaning Qiao, Andrew Dawson, Tony Parry, and Gerardo W. Flintsch. “Immediate effects of some corrective maintenance interventions on flexible pavements”. In: *Int. J. Pavement Eng.* 19.6 (2018), pp. 502–508.
[20] Samer Madanat. “Optimal infrastructure management decisions under uncertainty”. In: *Transp. Res. Part C* 1.1 (1993), pp. 77–88.
[21] C. P. Andriotis and K. G. Papakonstantinou. “Managing engineering systems with large state and action spaces through deep reinforcement learning”. In: *Reliab. Eng. Syst. Saf.* 191.April (2019), p. 106483.
[22] Konstantinos G. Papakonstantinou, Charalampos P. Andriotis, and Masanobu Shinozuka. “POMDP and MOMDP solutions for structural life-cycle cost minimization under partial and mixed observability”. In: *Struct. Infrastruct. Eng.* 14.7 (2018), pp. 869–882.
[23] Sheldon Ross. *Introduction to Probability Models*. 11th ed. Elsevier, 2014.
[24] Yi Jiang and Kumares C. Sinha. “Bridge service life prediction model using the Markov chain”. In: *Transp. Res. Rec.* 1223 (1989), pp. 24–30.
[25] Kamalesh Panthi. “A Methodological Framework for Modeling Pavement Maintenance Costs for Projects with Performance-based Contracts”. In: (2009), pp. 1–191.
[26] Ángela Alonso-Solorzano, Heriberto Pérez-Acebo, Daniel J. Findley, and Hernán Gonzalo-Orden. “Transition probability matrices for pavement deterioration modelling with variable duty cycle times”. In: *Int. J. Pavement Eng.* 24.2 (2023).
[27] Yasunari Inamura. “Estimating Continuous Time Transition Matrices From Discretely Observed Data”. In: *Bank Japan Work. Pap. Ser.* E07.6 (2006), pp. 1–40.
[28] Ward Whitt. “Continuous-time Markov chains”. In: *Dept. of Industrial Engineering and Operations Research, Columbia University, New York* (2006), p. 164.
[29] Toualith Jean Marc Meango and Mohamed Salah Ouali. “Failure interaction model based on extreme shock and Markov processes”. In: *Reliab. Eng. Syst. Saf.* 197 (2020).