Managing aging bridges under seismic hazards through deep reinforcement learning

Metwally, Z.; Andriotis, C.P.; Molaioni, F.

DOI
10.1201/9781003483755-403

Publication date
2024

Document Version
Final published version

Published in
Bridge Maintenance, Safety, Management, Digitalization and Sustainability

*Bridge Maintenance, Safety, Management, Digitalization and Sustainability –
Jensen, Frangopol & Schmidt (eds)*

Managing aging bridges under seismic hazards through deep reinforcement learning

Z. Metwally & C.P. Andriotis
Faculty of Architecture and the Build Environment, Delft University of Technology, Delft, The Netherlands

F. Molaioni
Department of Civil Engineering and Computer Science Engineering, University of Rome Tor Vergata, Rome, Italy

ABSTRACT: Structural systems must satisfy multiple performance and functionality requirements during their life cycle, withstanding safety-reducing degradation mechanisms and hazards. Intervention strategies must be planned accordingly to maintain structural integrity and minimize total life-cycle costs and risks, posing a complex optimization problem. Recent advances in multi-agent deep reinforcement learning (DRL) in conjunction with partially observable Markov Decision Processes (POMDPs) have shown great potential for determining optimal structural integrity management policies for systems with large state and action spaces compared to traditional decision practices. This paper tackles the maintenance optimization problem of aging bridges in seismic-prone areas, creating an updatable environment that embeds chloride-induced corrosion and state-dependent seismic fragility throughout the bridge life-cycle. The evolution of the environment is captured by a dynamic Bayesian network, and it is further integrated with decentralized multi-agent DRL algorithms to identify near-optimal life-cycle decisions under risk constraints. Results on a multi-component bridge system show the suitability of the developed framework for minimizing expected life-cycle costs, and for providing detailed and adaptive policies that significantly outperform traditional condition- and time-based maintenance plans.

1 INTRODUCTION

Structural degradation of bridges, induced by time-dependent stressors and natural hazards, is a major threat to long-term structural safety. Controlling it becomes more and more crucial in the expanding and complexifying landscape of infrastructure networks. Many highway bridges worldwide have been constructed with outdated design principles, and are experiencing structural weakening due to aging and deterioration (Shekhar & Ghosh 2021). Late reports show, for instance, that almost half of the US bridges are more than 50 years old, with 35% requiring maintenance interventions of circa $125B. This raises the risk of disruptive transportation network closures, mobility restrictions, and potential failures (ASCE 2021). These figures highlight a broader challenge facing many countries: a growing stock of aging infrastructure that are reaching or have well exceeded their design life.

Among different deterioration mechanisms, corrosion, which is typically initiated by the penetration of aggressive substances (e.g., chlorides), is a threat that can compromise the ability of structural systems to withstand the design loads, resulting in a decreased structural reliability and, hence, increased vulnerability to natural hazards (Choe et al.

DOI: 10.1201/9781003483755-403

2008). Specifically, corrosion of reinforcement bars, which is typically associated with reinforcement mass loss, spalling of the concrete cover, and bond slippage, is shown to affect the strength and ductility of bridge components such as RC columns, bridge piers, and bearings (Rinaldi et al. 2022). The phenomenon manifests itself through uniform area reduction along the rebar length, as well as nonuniform spatially distributed local pits or cavities (Bertolini et al. 2004).

It is well understood that the severity of seismic consequences on RC bridges is magnified due to corrosion. Therefore, recent studies have developed formulations to quantify seismic fragility under corrosion. A common assumption is to consider uniform corrosion, therefore, uniform cross-sectional area losses (Alipour et al. 2011; Rao et al. 2017; Deng et al. 2018). Yet in more recent studies, it is shown that pitting corrosion can be more critical compared to uniform corrosion, especially in RC bridge columns, as it can reduce the lateral load capacity due to localized weakening (Shekhar & Ghosh 2021). Accounting for both effects, Molaioni et al. 2023 proposed a DBN approach to continuously track the seismic fragility of reinforced concrete bridges throughout the life-cycle. Bridges are treated as multi-component systems that include columns and different types of bearings. The aging of bridges is rigorously incorporated in the approach, considering mass loss due to pitting corrosion and associated secondary effects of strength reduction, cracking, and loss of confinement, in addition to deterioration of steel bearings and corresponding modified cyclic behavior. Overall, literature suggests that modeling aging bridges as heterogeneous multi-component systems, and considering uniform and pitting corrosion alike for concrete members, in addition to reduction of strength and ductility of rebars, loss of concrete cover, and deterioration of steel fixed and expansion bearings, can significantly enhance our predictive capabilities. The next question to be addressed is how we can harness such predictive models to synthesize precise optimal life-cycle maintenance and retrofit plans given the complexity and uncertainties of the involved systems and phenomena.

Over the past years, there has been increasing attention to the development of life-cycle decision optimization methodologies for infrastructure management, aimed at meeting this challenge in the most eco-socio-economically aware ways (Arunraj & Maiti 2007; Bucher & Frangopol 2006). Conventional approaches for determining maintenance policies for structural systems build largely upon heuristic decision rules, which are based on engineering judgment about intervals or thresholds for interventions and observations. Such plans can be optimized either through exhaustive search or genetic algorithms. Another family of advanced optimization formulations sets up the problem as a Partially Observable Markov Decision Process (POMDP), solving the sequential decision-making task in a closed-loop fashion by incorporating observations and uncertainties of the underlying environment (Papakonstantinou & Shinozuka, 2014). Solutions can be traced effectively through value-based iteration methods in low- to medium-sized systems (Papakonstantinou & Shinozuka, 2014; Papakonstantinou et al., 2018), or through Deep Reinforcement Learning (DRL) in large-scale ones (Andriotis & Papakonstantinou 2019, 2021).

Embracing these advances, we articulate the seismic management problem of aging bridges as a POMDP and solve it through Deep Decentralized Multi-agent Actor Critic (DDMAC) DRL architectures (Andriotis & Papakonstantinou, 2021). DDMAC allows us to deal with the curse of dimensionality, namely the exponential explosion of state and action spaces with an increasing number of components. The actors assimilate centralized system information and produce a decentralized policy for each component, guided by a single critic network which is a surrogate of the life-cycle cost. Thereby, DDMAC allows actor outputs to scale linearly with components, while exploiting system-level information. Constraints are also incorporated either through state augmentation or a primal-dual optimization approach. A sparser version of DDMAC relying on local component information and centralized training with decentralized execution is introduced in (Saifullah et al. 2024). The temporal behavior of the structural system and its stressors are described by a Dynamic Bayesian Network (DBN), as developed by (Molaioni et al. 2023). This allows uncertainty propagation, state updating due to observations, and coupling with DRL in a principled manner (Morato et al. 2023). The environment states are defined as probabilistic distributions, i.e., beliefs, rather than deterministic scalars.

Overall, this paper introduces a framework that integrates chloride-induced deterioration of components, state-dependent seismic fragility, probabilistic hazard analysis, maintenance actions, associated costs, and structural risk with DRL to optimize life-cycle intervention policies. Exemplifying the approach in a case study of pre-70's bridges, structures are idealized as multi-component systems that include columns, high-type expansion bearings, high-type fixed bearings, and low-type fixed bearings. The aging of bridges is rigorously incorporated considering mass loss due to uniform and pitting corrosion, and associated secondary effects (Molaioni et al. 2023). Our approach results in substantial improvement in minimizing the life-cycle costs compared to traditional maintenance policies, offering detailed and dynamic maintenance plans at the component level, which are, however, risk-compliant and near-optimal at the system-level.

2 SEQUENTIAL DECISION-MAKING

The sequential decision-making problem is formulated as a Partially Observable Markov Decision Process (POMDP). Transfer between states is defined by a transition probability model, which is the probability of moving to a new state, s_{t+1}, in the coming time step, given the current state, s_t, when taking action a_t, i.e., $\Pr(s_{t+1} \mid s_t, a_t)$. Based on s_{t+1}, s_t, and the imposed action a_t, a cost c_t is received, which is a measure of the desirability of the environment state and imposed action. The environment is perceived through noisy observations, o_t, conditional on the actual hidden state, s_t, and the imposed action a_{t-1}, according to an observation model which defines the probability of observing o_t given s_t and a_t, i.e., $\Pr(o_t \mid s_t, a_{t-1})$. Leveraging the transition and observation models, sufficient statistics that encode all the previous observation and action history, i.e., beliefs, can be formed which can be forwarded to the agent. The agent is tasked with learning the best mapping between belief b_t and potential actions, i.e., a policy $\pi=\Pr(a_t \mid b_t)$.

Throughout training, π is constantly updated to determine an optimal policy π^* within a region of feasible policies Π_c, defined by the set of stochastic and deterministic constraints of the optimization problem (Andriotis and Papakonstantinou, 2021):

$$\pi^* = \underset{\pi \in \Pi_c}{\arg\min}\, \mathbb{E}\left(\sum_{t=0}^{T} \gamma^t c_t |\ a_t \sim \pi(o_{0:t},\ a_{0:t-1}), b_0 \right) \tag{1}$$

where γ is a positive discount factor lower than 1.0 translating future costs to the current value.

The optimized policy π^* can be written in terms of the initial belief b_0, which is updatable after each decision step in a Bayesian manner in the POMDP framework, as follows:

$$\pi^* = \underset{\pi \in \Pi_c}{\arg\min}\, V^\pi(b_0) \tag{2}$$

where V^π is the value function, which is a surrogate of the expected cumulative costs throughout the life-cycle until termination, $t=T$.

Here, the actor and critic neural networks parameterize the policy and value functions, respectively, following an off-policy approach, as follows:

$$g_{\theta}{}^i = \mathbb{E}\left[w_t \nabla_{\theta}{}^i \log \pi^i\left(a_t^i | b_t; \theta^i\right) A^\pi(b_t, a_t)\right] \tag{3}$$

$$g_\varphi = \mathbb{E}\left[w_t \nabla_\varphi V^\pi(b_t) A^\pi(b_t, a_t)\right] \tag{4}$$

where i is the index of the i^{th} actor; θ^i is the actor neural network weights vector; φ is the critic neural network weights vector;, w_t is an importance sampling weight; and A^π is the advantage function. The advantage function used takes the form:

$$A^\pi(b_t, a_t) = c(b_t, a_t) + \gamma V^\pi(b_{t+1}; \varphi) - V^\pi(b_t; \varphi) \tag{5}$$

3 STATE-DEPENDENT SEISMIC FRAGILITY

To quantify the seismic risk, we need to integrate the probability of exceedance of specific damage states due to distinct seismic intensity measures, over the range of these intensity measures, based on the bridge site characteristics:

$$PA_s(t) = \int \Pr[SDS(t) > s_s | IM_t = x] \left| \frac{dH(x)}{dx} \right| dx \tag{6}$$

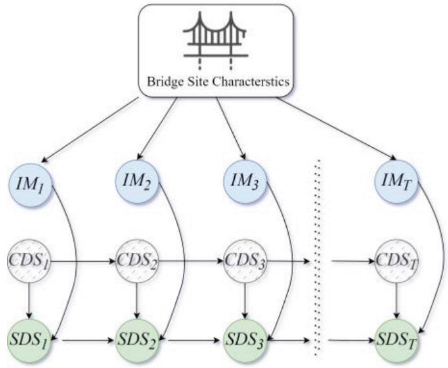

Figure 1. DBN for seismic fragility of aging bridges throughout the life-cycle.

where $PA_s(t)$ refers to the annual probability of exceedance of the seismic damage state (SDS=s_s) at time t; $H(x)$ is the hazard curve that quantifies the annual probability of exceedance given intensity measure IM_t=x, where x is typically the PGA magnitude at the bridge site.

To consider the effect of aging on the seismic fragility, it is important to consider the historical deterioration of the bridge component due to environmental factors (e.g., chloride-induced corrosion). In order to take this into account, the deterioration due to corrosion is modeled using non-stationary Markovian transition models conditioned on the prior corrosion state at time t-1, Pr(CDS_t=s_c| CDS_{t-1}). Following the same rationale, the probability of exceedance of the seismic damage state, SDS_t = s_s, at time t, is defined given the associated corrosion damage state, the intensity measure, and the prior seismic damage state, Pr(SDS_t ≥ s_s | CDS_t, IM_t, SDS_{t-1}). The corresponding DBN is graphically shown in Figure 1. Accordingly, at a specific time, t, of the bridge life, associated with a nonstationary deterioration rate, the seismic fragility, given the complete histories of corrosion deterioration, seismic damages, and prior seismic hazards, is approximated as follows:

$$\Pr(SDS_t \geq s_s | CDS_{0:t}, IM_{0:t}, SDS_{0:t}); \Pr(SDS_t \geq s_s | CDS_t, IM_t, SDS_{t-1}) \tag{7}$$

4 APPLICATION

The developed framework, integrating state-dependent seismic fragility assessment and life-cycle decision optimization through multi-agent DRL, is applied to a case study of a transportation corridor comprising two bridges in series, as shown in Figure 2. Note that spatial ground motion variability is not examined here, i.e., the accelerograms for the two bridges are identical. The bridges are located in a seismic-prone area (Monti et al. 2023), with a hazard curve shown in Figure 2. Each bridge is a four-component system that contains columns (COL), High-Type Expansion Bearings (HTEB), High-Type Fixed Bearings (HTFB), and Low-Type Fixed Bearings (LTFB). The corrosion of the bridge components is defined in a discrete space that contains four CDSs, i.e., Sound, Initial, Progressive, Critical. The corrosion state of the columns (COL) is discretized according to the percentage of rebar mass loss, while the discretization for the bearings depends on the mass loss of the bolts, the reduction in steel plate thickness, and the additive coefficient of friction due to corrosion-induced interlocking effects. Readers with more interest in the modeling details are referred to (Molaioni et al. 2023). Similar to corrosion, the seismic damage states are discretized into five SDSs, i.e. Intact, Slight, Moderate, Extensive, Complete, with the latter corresponding to (near-)failure. Here, the five states are defined probabilistically using log-normal distributions for the curvature ductility and displacement limits for columns and bearings, respectively (Nielson, & DesRoches 2006). Correlations between component SDSs can be modeled as in (Song & Kang, 2009). However, they are omitted here in the interest of safety due to the lack of redundancy in the considered series system.

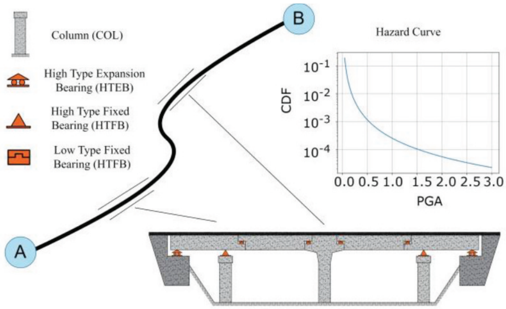

Figure 2. System information including bridge schematic and hazard curve.

Without loss of generality, seismic states are considered fully observable, whereas to capture the imprecision of annual corrosion state inspections, it is assumed that CDSs are observed with noise according to the following observation matrix (Andriotis & Papakonstantinou 2021):

$$\Pr(o_c|CDS_t = s_c) = \begin{array}{c} \\ s_{c,1} \\ s_{c,2} \\ s_{c,3} \\ s_{c,4} \end{array} \begin{array}{c} \begin{array}{cccc} o_{c,1} & o_{c,2} & o_{c,3} & o_{c,4} \end{array} \\ \begin{pmatrix} 0.84 & 0.13 & 0.02 & 0.01 \\ 0.11 & 0.77 & 0.09 & 0.03 \\ 0.02 & 0.16 & 0.70 & 0.12 \\ 0.01 & 0.02 & 0.13 & 0.84 \end{pmatrix} \end{array} \quad (8)$$

The aim is to optimize structural interventions over a 50-year planning horizon while constraining the expected risk of disconnection between points A and B to no greater than 5% in 50 years. Four maintenance actions are considered for columns (COL) and two actions for bearings (HTEB, HTFB,LTFB). These are shown in Table 1 in conjunction with the corresponding costs and effects on the corrosion and seismic damage states. In addition to these actions, the agent can choose to not interfere with the environment at a given decision step (i.e., Do Nothing). The costs associated with the potential actions are based on actual contract prices (Tennessee Department of Transportation 2023, Ghosh & Padgett 2011). An additional mobilization cost of 15% of replacement cost is also imposed, which incentivizes taking simultaneous actions for similar components. Note that identical components within the same bridge are treated as one component within this setting.

In order to optimize the life-cycle cost considering the aforementioned maintenance actions, costs, and risk constraints, the DDMAC and DDMAC-CTDE multi-agent algorithms are adopted for policy training. In both algorithms, the actor network consists of 8 separate control units, corresponding to the components of the two-bridges corridor, with no parameter sharing among them. The two bridges work in series, i.e., losing one causes disconnection between A and B. The control units corresponding to the columns (COL) have a 5-dimensional softmax output, while the rest have a 3-dimensional output, which corresponds to the number of potential intervention actions. DDMAC control units share the system information, i.e., each unit has access to other units' information in addition to its

Table 1. Potential actions for each component, corrosponding cost, and state effects.

Component	Action	Cost	Action Effect
Columns (COL)	Thin Epoxy overlay	2149$	Restart deterioration rate;
	Epoxy Injection	3207$	Improve CDS & SDS by one; Delay CDS deterioration rate by 5 years;
	Concrete Repairs	2388$	Improve CDS & SDS by two;
	Replace	20000$	Bring to intact condition
Bearings (HTEB, HTFB, LTFB)	Replace Anchor Bolts	1112$	Improve CDS & SDS by one;
	Replace	2397$	Transition to intact condition

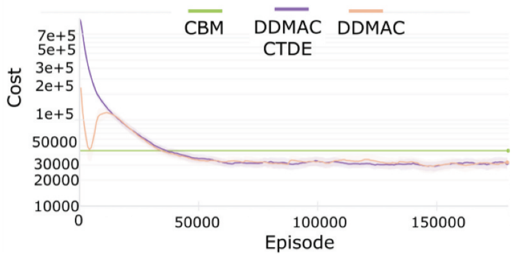

Figure 3. Training performance of DRL algorithms and comparison with baseline policy.

own. On the contrary, DDMAC-CTDE units adopt a decentralized information approach, where each unit accesses only its own information. In both algorithms, the actor units are guided by the critic network, which approximates the value function of the entire system with a one-dimensional linear output. For evaluating the performance of the learned policy, the DRL-based policy is compared to a condition-based baseline (CBM), where a corrective policy typically adopted to mitigate seismic damage (Ghosh & Padgett 2011) is imposed on top of a heuristically optimized condition-based maintenance policy on CDSs (CBM). Three training instances of each of the adopted DRL algorithms trained for 180K episodes, along with the optimized CBM, are shown in Figure 3. Both DDMAC and DDMAC-CTDE produce comparable training averages, and both overperform the optimized CBM by almost 28% in terms of expected life-cycle costs. It was observed that the trained policies managed to also adhere to the risk constraints, with a tendency to prescribe more preventive action in relatively early life-cycle phases, as shown by the cumulative cost and probability of failure evolution of 1000 random realizations shown in Figure 4.

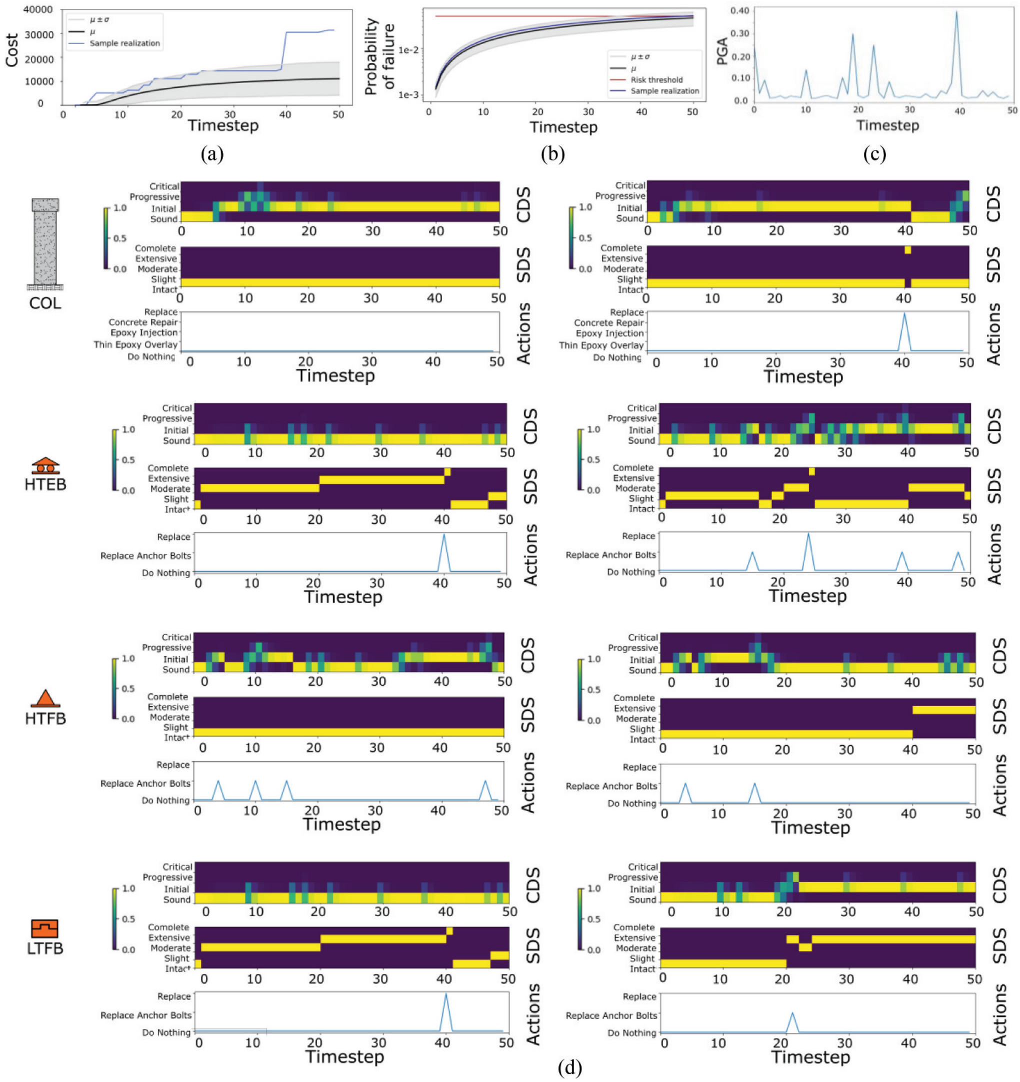

Figure 4. Policy realization: (a) Cost accumulation; (b) Probability of failure evolution; (c) Life-cycle seismic shocks; (d) CDS beliefs, SDSs, and actions.

To visualize the dynamics of the trained policy, a sample realization is presented in Figure 4, showing the observed seismic shocks characterized by the corresponding PGA on a temporal scale, followed by the corrosion CDS belief contours, the observed seismic SDS, and the imposed actions for each component. First, the trajectory starts with observing a relatively intense seismic event, with PGA≈0.25. This results in no seismic damage, except for damage in the HTEB components of both bridges, which are associated with the height fragility (Molaioni et al. 2023), in addition to slight damage in LTFB-1. While the agent decides to take no reactive action upon the damage, an intervention is made when corrosion deterioration is observed in other instances (e.g., HTEB-2 at year 15, LTFB-1 at year 4), which could have increased the vulnerability of such components to future shocks. Another major shock is observed around year 19, which causes different extents of damage across the bearing components, however, due to the prior preventive actions, no failure is observed throughout the system and the connectivity was maintained until two subsequent major shocks are observed at years 23 and 39. The latter causes failures in different spots, resulting in corrective replacement actions. It should be noted that the agent typically avoids maintaining the columns due to the relatively high associated costs, and due to their higher reliability compared to other components, as shown in detail in (Molaioni et al. 2023), except when failed. Similarly, the agent prefers to impose preventive repairs instead of costly replacements for bearings. Finally, the agent leverages the mobilization incentive, by imposing simultaneous interventions to HTFB-1 and HTFB-2 at years 4 and 15.

5 CONCLUSIONS

In this paper, we introduce a framework for managing aging bridges subject to seismic hazards using dynamic Bayesian networks (DBNs) and Deep Reinforcement Learning (DRL). A virtual environment incorporates chloride-induced corrosion of bridge components and state-dependent seismic fragility functions, which are embedded in a DBN to allow for uncertainty propagation and state information updating throughout the life-cycle. Leveraging recent advances in multi-agent DRL in conjunction with partially observable Markov decision processes, the developed environment is probed by DRL agents operating on the belief space of the system, i.e., the dynamically updatable probabilistic distributions over component states. The framework is applied to identify a near-optimal 50-year intervention policy for a two-bridge transportation corridor located in a seismic-prone area, as well as the corresponding expected life-cycle cost under long-term risk constraints. A substantial improvement of 28% is observed in minimizing the life-cycle costs with DRL compared to traditional maintenance policies, while adhering to the imposed risk constraints. The introduced framework can be adopted, without loss of generality, to the management of other hazard- and deterioration-vulnerable systems, enhancing the safety, reliability, and cost-effectiveness of critical infrastructure networks.

ACKNOWLEDGEMENTS

This material is based upon work supported by the TU Delft AI Labs Program. The authors gratefully acknowledge this support.

REFERENCES

Alipour, A., Shafei, B., & Shinozuka, M. (2011). Performance Evaluation of Deteriorating Highway Bridges Located in High Seismic Areas. *Journal of Bridge Engineering*, 16, 597–611.

Andriotis, C. P., & Papakonstantinou, K. G. (2019). Managing engineering systems with large state and action spaces through deep reinforcement learning. *Reliability Engineering & System Safety*, 191.

Andriotis, C. P., & Papakonstantinou, K. G. (2021). Deep reinforcement learning driven inspection and maintenance planning under incomplete information and constraints. *Reliability Engineering & System Safety*, 212.

ASCE (2021). 2021 report card for America's infrastructure. Retrieved 15 September 2023. Available from https://infrastructurereportcard.org/wp-content/uploads/2020/12/National_IRC_2021-report-2.pdf

Arunraj, N. S., & Maiti, J. (2007). Risk-based maintenance - Techniques and applications. *Journal of Hazardous Materials*, 142(3), 653–661.

Bertolini, L., Elsener, B., Pedeferri, P., Redaelli, E. and Polder, R.B., 2013. Corrosion of steel in concrete: prevention, diagnosis, repair. *John Wiley & Sons*.

Bucher, C., & Frangopol, D. M. (2006). Optimization of lifetime maintenance strategies for deteriorating structures considering probabilities of violating safety, condition, and cost thresholds. *Probabilistic Engineering Mechanics*, 21(1), 1–8.

Choe, D. E., Gardoni, P., Rosowsky, D., & Haukaas, T. (2008). Probabilistic capacity models and seismic fragility estimates for RC columns subject to corrosion. *Reliability Engineering & System Safety*

Deng, P., Zhang, C., Pei, S., & Jin, Z. (2018). Modeling the impact of corrosion on seismic performance of multi-span simply-supported bridges. *Construction and Building Materials*, 185 (Oct): 193–205.

Ghosh, J., & Padgett, J. E. (2011). Probabilistic seismic loss assessment of aging bridges using a component-level cost estimation approach. *Earthquake Engineering & Structural Dynamics*, 40(15), 1743–1761.

Molaioni, F., Rinaldi, Z., & Andriotis, C. P. (2023). Assessing life-cycle seismic fragility of corroding reinforced concrete bridges through dynamic Bayesian networks. In Life-Cycle of Structures and Infrastructure Systems (pp. 523–530). CRC Press.

Monti, G., Demartino, C., & Gardoni, P. (2023). Towards risk-targeted seismic hazard models for Europe. *Scientific reports*, 13(1), 10717.

Morato PG, Andriotis CP, Papakonstantinou KG, Rigo P. (2023). Inference and dynamic decision-making for deteriorating systems with probabilistic dependencies through Bayesian networks and deep reinforcement learning. *Reliability Engineering & System Safety*.

Nielson, B.G. and DesRoches, R., 2006. Seismic fragility methodology for highway bridges. In *Structures Congress 2006: Structural Engineering and Public Safety* (pp. 1–9).

Papakonstantinou, K. G., & Shinozuka, M. (2014). Planning structural inspection and maintenance policies via dynamic programming and Markov processes. Part II: POMDP implementation. *Reliability Engineering & System Safety*.

Papakonstantinou, K. G., Andriotis, C. P., and Shinozuka, M. (2018). POMDP and MOMDP solutions for structural life-cycle cost minimization under partial and mixed observability. *Structure and Infrastructure Engineering*, 14(7), 869–882.

Rao, A. S., Lepech, M. D., Kiremidjian, A. S., & Sun, X.-Y. (2017). Simplified structural deterioration model for reinforced concrete bridge piers under cyclic loading. *Structural Infrastructure Engineering*, 13 (1): 55–66.

Rinaldi, Z., Di Carlo, F., Spagnuolo, S., & Meda, A. (2022). Influence of localized corrosion on the cyclic response of reinforced concrete columns. *Engineering Structures*.

Saifullah, M., Papakonstantinou, K. G., Andriotis, C. P., & Stoffels, S. M. (2024). Multi-agent deep reinforcement learning with centralized training and decentralized execution for transportation infrastructure management. arXiv preprint arXiv:2401.12455.

Shekhar, S., & Ghosh, J. (2021). Improved component-level deterioration modeling and capacity estimation for seismic fragility assessment of highway bridges. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*, 7(4), 04021053.

Song, J. and Kang, W.H., 2009. System reliability and sensitivity under statistical dependence by matrix-based system reliability method. *Structural Safety*, 31(2), pp.148–156.

Tennessee State Government. (2023). Retrieved 01 November 2023. Available from https://www.tn.gov/

Zhang, N. L., & Si, W. J. (2020). Deep reinforcement learning for condition-based maintenance planning of multi-component systems under dependent competing risks. *Reliability Engineering & System Safety*, 203.