Observation-based Decision-making for Infrastructure

Eleni Chatzi[1], Konstantinos G. Papakonstantinou[2], Rade Hajdin[3], Daniel Straub[4]

DOI: https://doi.org/10.5592/CO/BSHM2017.4.1

[1]Associate Professor, Department of Civil, Environmental and Geomatic Engineering, ETH Zurich, Stefano-Franscini-Platz 5, 8093, Zurich, Switzerland
[2]Assistant Professor, Department of Civil & Environmental Engineering,The Pennsylvania State University, University Park, PA 16802-1408, USA
[3]Professor, University of Belgrade, Faculty of Civil Engineering, Kralja Aleksandra 73, 11000 Belgrade, Serbia
[4]Associate Professor, TUM Department of Civil, Geo and Environmental Engineering, Technical University of Munich, Arcisstrasse 21, 80290 München, Germany

E-mails: [1]chatzi@ibk.baug.ethz.ch; [2] kpapakon@psu.edu; [3]rade.hajdin@grf.bg.ac.rs; [4]straub@tum.de

Abstract. The aim of this review paper is to summarize available options for decision-support toward the optimal management of infrastructure systems under consideration of uncertainties. Thereafter, we elaborate on a particular variant, namely Partially Observable Markov Decision Processes (POMDP), as a flexible framework for the incorporation of information from observations (visual inspections, non-destructive testing and monitoring) in the context of optimal inspection and maintenance planning. Examples from the literature are presented, demonstrating the potential of planning under uncertainty and the links to the Value of Information (VoI) from Structural Health Monitoring (SHM).

Keywords: Optimal Inspection and Maintenance Planning, Structural Health Monitoring (SHM), Value of Information (VoI), Markov Decision Processes (MDPs), Partially Observable Markov Decision Processes (POMDPs)

1 Introduction

Structures and infrastructure systems face challenges due to aging, deterioration and adverse operational conditions. Recent technological advances have allowed for development of Structural Health Monitoring (SHM) systems, which provide information on the “health” state of structural systems, and may be exploited to derive indicators of corresponding performance. SHM systems may therefore serve for supporting decisions regarding the management of infrastructure systems throughout their life-cycle. Such decisions pertain to the planning of appropriate inspection and maintenance actions in evidence of damage or deterioration. To address these challenges effectively and scientifically, new methods and tools are needed to quantify and optimize the Value of Information (VoI) from the SHM systems.

As part of COST Action TU1402, Working Group 3 aims at identifying developing and critically overviewing methods and tools required for the utilization of the theoretical VoI framework in infrastructure practice. Such methods take basis in modern methods of probabilistic systems analysis including Fully and Partially Observable Markov Decision Processes, Semi Markov Decision Processes, Bayesian Networks, Monte Carlo simulation schemes, Stochastic Meshing Algorithms, First Order Reliability Methods, and combinations thereof. The Influence Diagram offered in Fig. 1 illustrates the separate components involved in the Value of Information framework within the context of SHM (Straub et al., 2017).

The goal of such a framework is to provide support for optimal maintenance and intervention planning. Consider the example of a bridge object, where the goal lies in assurance of a desired service quality with minimum interruptions. In such a case, bridge owners launch preventive actions when the risk of service impairment, interruption or losses in life cycle costs reaches some predefined level. Implicitly the owners define the accepted risk based on socio-economic equity principles. This accepted risk depends upon the established performance goals for each component or combination of bridge components and together with the costs implied for every action (in the form of inspection or intervention) governs the policy to be followed for management of the system. In this context, as part of COST Action TU1406, Working Group 3 aims in defining the steps required for setting up Quality Control (QC) plans for diverse types of roadway bridges.

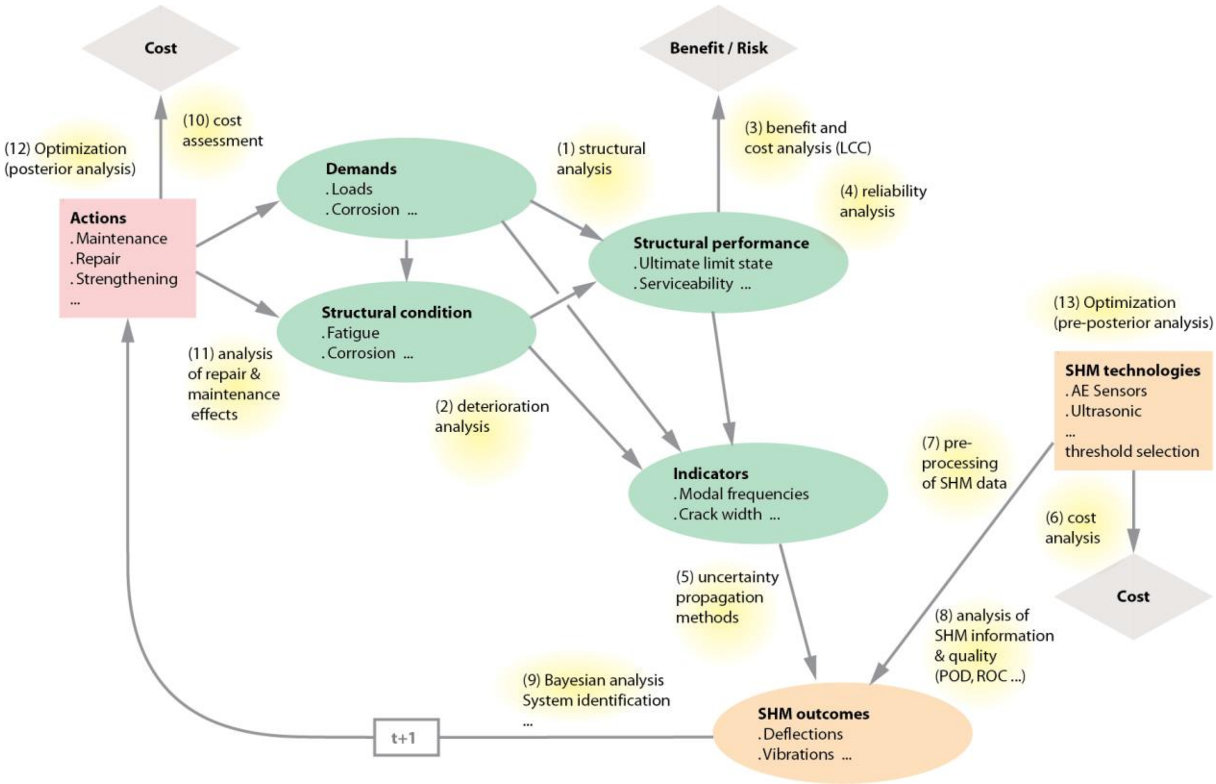

Fig. 1. Influence diagram (ID) of the Value of Information for SHM. Structural parameters and models are indicated in green, SHM parameters and models are denoted in orange, while repair, maintenance and related actions are marked in red. The yellow bubbles are the analysis methods and tools used in the different parts of the process. The box [t+1] indicates that the edge is from one time step to the next, hence this ID represents a decision process in time. Figure from Straub et al. (2017).

In existing literature, several approaches have been formulated to tackle the previously described problem of decision-making for infrastructure. A first take lies in casting this in an optimization framework, where the decision maker may choose to either optimize for separate objectives, usually tied to corresponding performance indicators, e.g. condition, availability, safety, or durability (Liu et al., 1997; Miyamoto et al., 2000; Furuta et al., 2004), or instead decide to simultaneously treat conflicting objectives. When adopting a Multi-Criteria Decision-Making (MCDM) approach, the preferred policy structure of the decision maker is adopted to transform the multiple objectives into a single optimization function. Approaches of this class, as summarized in Adiel Teixeira de Almeida et al. (2015), include single criterion-synthesis methods such as the Multi-attribute Utility Theory and the Multi-attribute Value Theory (Keeney & Raiffa, 1976), outranking methods such as Elimination and Choice Expressing Reality (ELECTRE) described in Figueira et al. (2005) and the Preference Ranking Organization Method for Enrichment of Evaluations (PROMETHEE) overviewed by Brans & Mareschal (2005), as well as further alternatives. As another option, multi-objective optimization may be employed to tackle multiple objectives, delivering a set of compromising decision options along the so-called Pareto front. To this end, Liu and Frangopol (2005) employ multi-objective optimization to solve the combinatorial optimization problem of annual prioritization of maintenance efforts for deteriorating components of concrete bridges. Taflanidis & Beck (2008) proposed a robust stochastic design framework in the context of reliability analysis and related decision support, where probabilistic models of excitation uncertainties and system modeling uncertainties can be introduced. The two-stage framework implements Stochastic Subset Optimization (SSO) for identifying a region of interest in the design space, and a stochastic optimization algorithm to finally identify the optimal solution.

An issue that inevitably enters the decision-making process is that of uncertainties (aleatory or epistemic). In this respect engineering decision problems may be classified into three main categories, namely those of prior, posterior, or pre-posterior decision problems, as elaborated upon by Faber (2005). Pure prior and posterior decision problems, as for instance calibration of code format or service life extension, may be solved using standard approaches of Quantitative Risk Analysis (QRA) and Structural Reliability Analysis (SRA). However, efficient treatment of the broader pre-posterior decision class, which includes the problem of inspection and maintenance planning, necessitates dedicated tools (Faber, 2003; Sørensen, 2009), such as Bayesian Probabilistic Networks and Influence Diagrams (Heckerman, 1995). In this work, we elaborate on a candidate approach, namely Partially Observable Markov Decision Processes (POMDPs). Markov Decision Processes (MDPs) have been widely adopted across diverse domains of engineering and applied sciences in the context of decision-making. A POMDP is in essence a generalization of a Markov Decision Process (MDP), where the decision-maker does not hold complete, i.e., deterministic, information on the system state (condition), as well as on the outcome of the performed inspections/actions.

2 An MDP approach to Decision-making

The motivation behind utilization of the POMDP framework lies in its capability to (i) incorporate stochastic models and uncertain data based on firm mathematical foundations, (ii) to optimize across long-term objectives, and (iii) to incorporate near-real-time observations allowing for near-real-time optimal decision support.

2.1 Markov Decision Processes (MDPs)

A MDP comprises the following basic components: i) the states, which in the context of infrastructure may be tied to indicators of the system's performance or condition (Limongelli et al., 2005); ii) a set of actions, commonly pertaining to interventions (repair, retrofitting, replacement); iii) the consequences of actions on the system's state; and iv) the value/cost of these actions (Cassandra et al., 2005). The MDPs abide to the Markov assumption, implying that the state at a given decision (time) step is uniquely determined by the previous state. Various solution methods have been developed for MDP problems, with Value Iteration (Bellman, 1957) and Policy Iteration (Howard,1960) algorithms holding the lead.

MDPs have long served as tools for decision support. As an example of such a use-case, the Swiss Federal Roads Office employs MDPs for bridge management, as part of the KUBA software (Hajdin, 2008). In KUBA-MS, the assessment units are structural elements, which may be further divided into segments. An influence indicator is linked to each segment, according to its exposure to environmental influences. Prevalent deterioration processes are identified and resulting states are rated according to a five-scale classification system, ranging from 1 (good condition, i.e., no damages) to 5 (alarming condition, i.e., urgent actions necessary). Markov chains are used to obtain condition forecasts and MDP use used to determine optimum preservation policies (Fig. 2). The transition matrices result from statistical analysis of inspection data (Hajdin & Peeters, 2008).

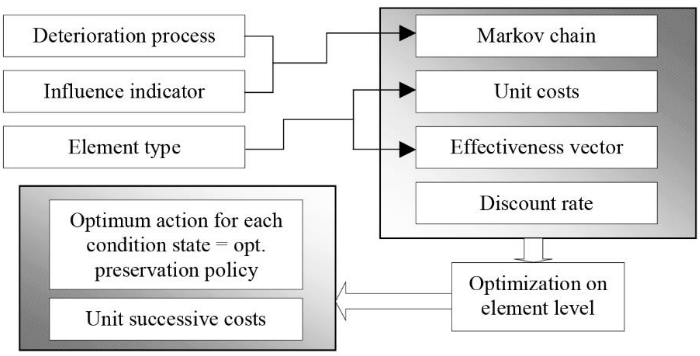

Fig. 2. KUBA-MS Optimization on the element level. Figure reused from (Hajdin, 2008).

2.2 Partially Observable Markov Decision Processes (POMDPs)

MDPs operate on the assumption of complete (deterministic) knowledge of the system's state, which is hardly ever the case. An educated estimate on the condition of the system may be delivered via inspection, non-destructive evaluation or monitoring technologies; none of these options however provide complete and exact information on the system's state. To address this challenge, Partially Observable Markov Decision Processes (POMDPs) relax the MDP assumption to consider probabilistic knowledge of the state, allowing for the analysis of decision-making under uncertain observations.

The POMDP describes the interaction of an agent, in this case a structure, with its surrounding world via the tuple $\{S, A, T, \Omega, O, R\}$:

S : A set of (usually discrete) system states.
A : A set of possible, usually discrete, control actions.
T : A transition model describes the evolution of the system from a state s_i at step i to a future state s_{i+1}, as a consequence of an executed action a. The system's state is updated using the conditional probability $p(s_{i+1} \mid s_i, a)$.
Ω : A set of (usually discrete) observations describing the outcome of an inspection or monitoring method.
O : The observation model defines the probability of obtaining an observation outcome o when the system lies in state s, i.e., $p(o \mid s)$.
R : A reward (or cost) function outputs a reward value as $r \in \Re$. The reward may be modeled as dependent upon the current state, the assumed action or a combination of both.

The POMDP process is summarized as indicated in Fig. 3. The agent initiates at a state s_i. An observation o_i is taken, motivating execution of action a_i, shifting the system's state to s_{i+1}, which results in a reward r_i. The state is only partially known due to the uncertainties involved in the system's observation, thus it may only be described in terms of a belief (a probability distribution) over the state-space. Once the system evolves due to the agent's action, the belief state is updated based on the previous belief state, the executed action, and the received observation. This is performed by employing Bayes' rule:

$$b^{a,o}(s_{i+1}) = \frac{p(o \mid s_{i+1})}{p(o \mid \mathbf{b}, a)} \sum_{s_i \in S} p(s_{i+1} \mid s_i, a) b(s_i) \tag{1}$$

where a, o designate the executed control action and inspection (observation), respectively, $b(s_i)$ is the current belief state describing our confidence in the system lying in state s_i.

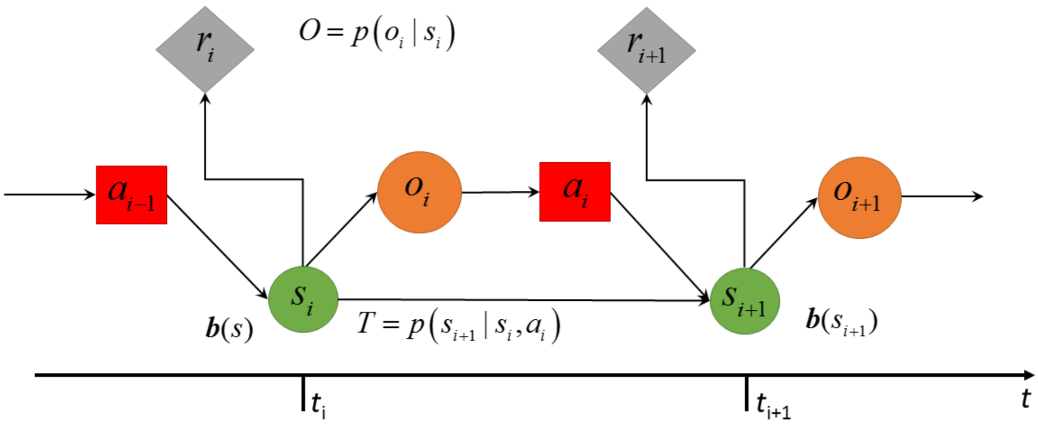

Fig. 3. The POMDP sequential decision process with alternating actions and inspections.

3 POMDP Solution Methods

Discrete POMDPs may be solved via a number of available solution algorithms (Poupart and Boutilier 2004, Sondik 1971). Most of these methods are approximate and rely on sampling of the belief space and propagation of these samples through the sequential planning process (observations, actions). Given the belief state $\boldsymbol{b}$ in a time horizon *n*, i.e., when *n* decision steps or decision intervals are left, and the updated belief state in horizon *n-1*, the optimal value function at horizon *n* is calculated in a recursive manner, as $V_n(\boldsymbol{b}) = \max_a Q_n(\boldsymbol{b}, a)$, where the expected reward Q_n for a given belief state $b(s)$ and action a is $Q_n(\boldsymbol{b}, a) = \sum_{s \in S} r_a(s) b(s) + \gamma \sum_o p(o \mid \boldsymbol{b}, a) V_{n-1}(\boldsymbol{b}^{a,o})$, where $\gamma \in [0,1)$ is the discount factor, which weighs the significance of the immediate in comparison to the delayed reward, $V_{n-1}(\boldsymbol{b}^{a,o})$ is the optimal value function of the previous horizon depending on the belief state $\boldsymbol{b}$, the current action a, and the observation o.

The goal of planning is to maximize the future expected reward by selecting an appropriate policy π^*, i.e., a sequence of actions and observations that maximizes the value function.

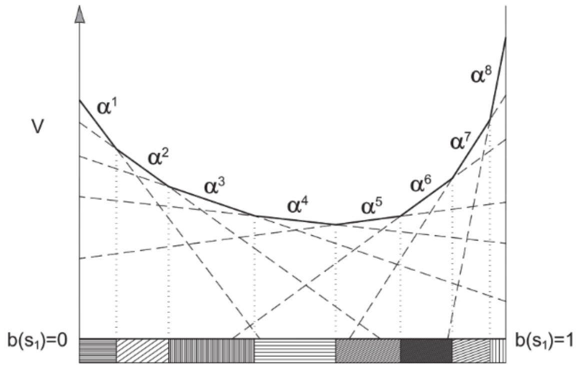

Fig. 4. Optimal a-vectors plotted in the belief space. Figure reused from Papakonstantinou & Shinozuka (2014a)

Sondik (1971) proved that the recursive calculation of the value function in a discrete system degenerates in the search for the so-called optimal a-vectors, each representing the value function for an optimal strategy at the current planning horizon n and for a set of belief states $b(s)$.

The belief space is a simplex, and each vector defines a region over the simplex, which represents a set of belief states. The value function is generally defined as the upper surface of these vectors.

The problem may be solved by means of Value iteration and Policy iteration algorithms; further alternatives are described in Papakonstantinou & Shinozuka (2014a). The POMDP problem may also be cast in the continuous space, as elaborated upon in the work of Schöbi & Chatzi (2016). This allows for more flexible formulations able to account for generic, nonlinear actions and observations. In this case, solution methods include policy search (Ng & Jordan, 2000), or grid- and point-based value iteration algorithms, that are extended to fit the continuous space (Porta et al., 2004).

4 POMDP Implementation in Infrastructure Management

POMDPs admittedly remain less popular than their MDP alternative in the domain of infrastructure planning and policy making, largely owing to their higher computational complexity. However, they offer numerous advantages, some of which are inherited from MDPs, such as flexibility in terms of formulation, functionality in both a discrete or continuous setting, inclusion of periodic and aperiodic inspection intervals, perfect and imperfect inspections, deterministic and probabilistic actions, stationary and non-stationary environments (also treatable within a semi-Markovian context), as well as planning in an infinite or finite decision horizon. In an early work, Madanat & Ben-Akiva (1994) adopt POMDPs for decision-making for highway-pavement networks. Ellis et al. (1995), and Corotis et al. (2005) demonstrate use of POMDPs for bridge inspection planning. Papakonstantinou & Shinozuka (2014a, 2014b) provide a thorough overview of solvers suited for solution of large-scale and more realistic problems.

In an example from (Papakonstantinou & Shinozuka, 2014b), a POMDP policy is described for decision support in the face of corrosion of the steel in a reinforced concrete wharf deck slab. A spatial stochastic corrosion model is used, as specified in (Papakonstantinou & Shinozuka, 2013), which defines the transition probability matrix. In accordance with AASHTO specifications, four discrete states (conditions) are assumed: condition 1 (less than 10% damage), condition 2 (damage between

10%- 25%), condition 3 (damage between 25%-50%), and finally condition 4 (over 50% damage). Implementation of the POMDP algorithm is exemplified in Fig. 5 for two consecutive steps of one possible state evolution scenario. In decision step 119 (top plot) the structure actually lies in condition 2, while the decision-maker attributes a mere 28% to this condition. The POMDP computed policy indicates a visual-inspection at the beginning of the next step. Once this is carried out, the belief is then updated to 81.20% with respect to the possibility that the system (indeed) lies in condition 2. For this belief state, the computed policy plan suggests execution of a minor repair action and subsequently indicates that visual inspection should be chosen as a cost-effective observation option, which will eventually witness the shift of the structure in the improved condition level 1.

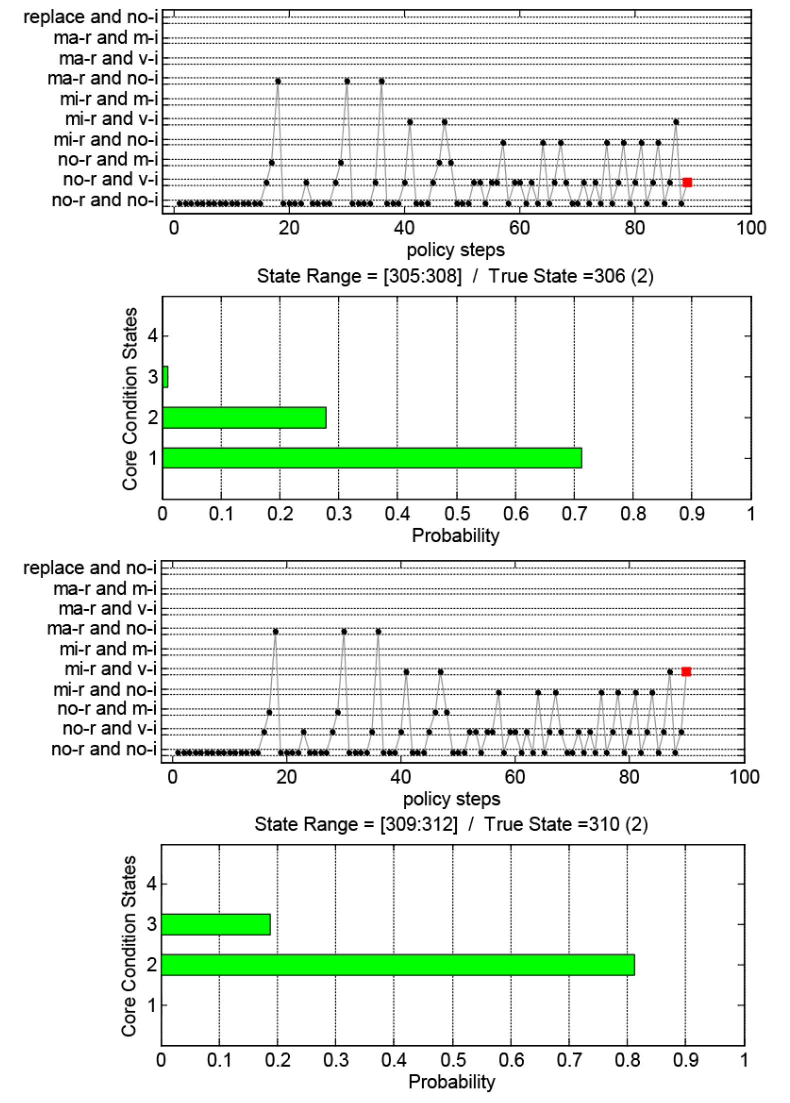

Available policy options:

replace and no-inspection
major-repair and monitoring-inspection
major-repair and visual-inspection
major-repair and no-inspection
minor-repair and monitoring-inspection
minor-repair and visual-inspection
minor-repair and no-inspection
no-repair and monitoring-inspection
no-repair and visual-inspection
no-repair and no-inspection

Fig 5. Detailed look at two consecutive policy steps (top to bottom) according to the POMDP computed policy. Figure reused from (Papakonstantinou & Shinozuka, 2014c).

The sequence of policy steps illustrated in Fig. 5 reflects a typical policy planning scenario for infrastructure management; the optimal policy typically favors major-repair actions in early stages, taking advantage of the improved and low deterioration rates. Instead, as demonstrated in a related example (Papakonstantinou et al., 2016), in the case where availability of permanent monitoring data is assumed, the algorithm will tend to exploit the more precise (in comparison to the visual-inspection) information provided by the monitoring system, and strategically suggest more frequent and inexpensive minor repair actions. The latter pro-actively prevents the evolution of severe damage, thus alleviating eventual need for replacement.

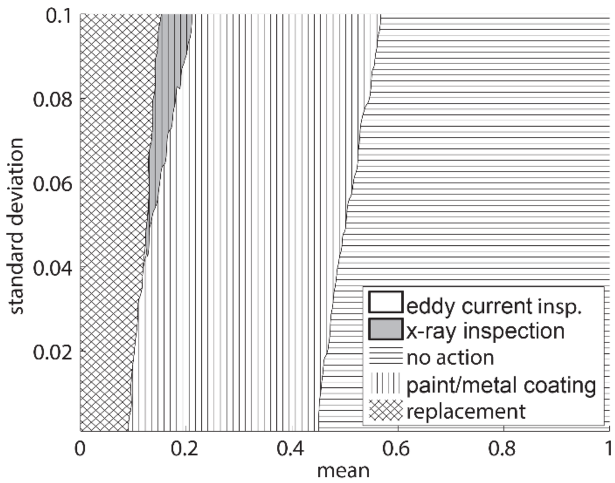

Fig. 6. Policy map for horizon 3 of the bridge maintenance problem in Schöbi & Chatzi (2016).

A valuable outcome of the POMDP policy planning process, lies in the possibility to produce decision maps, as illustrated for the continuous POMDP equivalent in Schöbi and Chatzi (2016). Once an indicator of the system's condition (state) is established, then the belief vector may be summarized via the first two moments (mean and variance).

Consequently, in a finite horizon problem setting, an optimal policy map may be formulated per horizon *n*, as illustrated in Fig. 6, indicating a recommended action (intervention) and corresponding inspection. Such policy maps may be exploited as decision support tools by operators and decision-makers.

POMDPs may further be extended to tackle problems of multiple components or systems, as exemplified in Memarzadeh & Pozzi (2016). The provided example on a farm of 25 wind turbine components reveals the importance of availability of observations from these components, in the form of inspections, for cost-effective management of the farm. However, POMDP are limited when considering the joint optimization of inspections and monitoring in a redundant structural system, due to computational costs. In this case, approximate solutions through heuristic policies may provide a pragmatic solution (Luque and Straub, under review).

5 Conclusions & Outlook

This review work outlines available alternatives for decision-making under uncertainty, shedding focus on the POMDP variant. This tool employs estimates on the condition of the system, extracted by means of uncertain or incomplete observations, for policy planning under the influence of deterioration processes, while assuming availability of actions whose effects are stochastic. Although not so far exploited in this sense, POMDPs and their computed optimal policy trajectories could be utilized in order to feed the components of pre-posterior analysis tools for quantifying the Value of Information of Structural Health Monitoring systems. This is to be explored as a next research direction by the authoring team.

Acknowledgements

The authors would like to acknowledge COST Action TU1402 on "Quantifying the Value of SHM" and COST Action TU1406 on "Quality specifications for roadway bridges, standardization at a European level" for making this work feasible through promotion of an extended network of collaborators across diverse domains.

References

Adiel Teixeira de Almeida, Rodrigo José Pires Ferreira and Cristiano Alexandre V. Cavalcante. A review of the use of multicriteria and multi-objective models in maintenance and reliability. *IMA J Management Math*, 26(3): 249-271, 2015. doi: 10.1093/imaman/dpv010

Richard Bellman. A Markovian Decision Process. *Journal of Mathematics and Mechanics*, 6(4):679-684, 1957.

Jean-Pierre Brans, Bertrand Mareschal. *Promethee methods. Multiple Criteria Decision Analysis: State of the Art Surveys* (Figueira J. Greco S. Ehrgott M. eds). New York: Springer, 2005.

Anthony Cassandra, Marian Nodine, Shilpa Bondale, Steve Ford and David Wells. Using POMDP-based state estimation to enhance agent system survivability. *IEEE 2nd Symposium on Multi-Agent Security and Survivability*, 11-20, 2005.

Eleni Chatzi and Konstantinos G. Papakonstantinou. A Partially Observable Markov Decision approach to decision-making. *Factsheet of COST TU1402: Quantifying the Value of Structural Health Monitoring*, 2017.

Hugh Ellis, Mingxiang Jiang and Ross B. Corotis. Inspection, maintenance, and repair with partial observability, *Journal of Infrastructure Systems*, 1(2):92-99, 1995.

Ross B. Corotis, Hugh Ellis and Mingxiang Jiang. Modeling of risk-based inspection, maintenance and life-cycle cost with partially observable Markov decision processes, *Structure and Infrastructure Engineering*, 1(1):75-84, 2005.

Michael H. Faber. On the Treatment of Uncertainties and Probabilities in Engineering Decision Analysis. *ASME. Journal of Offshore Mechanics and Arctic Engineering,* 127(3):243-248, 2005. doi:10.1115/1.1951776.

Michael H. Faber. New Approaches to Inspection Planning of Fatigue Damaged Offshore Platforms. *Journal of UK Safety and Reliability Society*, 23 (1):35–50, 2003.

José Figueira, Vincent Mousseau and Bernard Ro. Electre methods. *Multiple Criteria Decision Analysis: State of the Art Surveys* (Figueira J. Greco S. Ehrgott M. eds). New York: Springer, 2005.

Hitoshi Furuta, Takahiro Kameda, Yoshiko Fukuda and Dan M. Frangopol. Life-cycle cost analysis for infrastructure systems: life cycle cost vs. safety level vs. service life, in D. M. Frangopol, E. Brühwiler, M. H. Faber, and B. Adey, (eds.), *Life-Cycle Performance of Deteriorating Structures: Assessment, Design and Management*, ASCE, Reston, VA, 19–25, 2004.

Rade Hajdin. KUBA 4.0 – The Swiss Road Structure Management System. Transportation Research Board, *In proceedings: 10th International Bridge and Structure Management Conference*, Buffalo, New York, October 20-22, 2008.

David Heckerman. A Tutorial on Learning Bayesian Networks. Technical Report 95-06, Microsoft, 1995. Damage Indicators. Factsheet of COST Action TU1402, 2015.

Rare Hajdin and Leon Peeters. Bridging Data Voids: Advanced Statistical Methods for Bridge Management in KUBA. Transportation Research Board, *In proceedings 10th International Bridge and Structure Management Conference*, Buffalo, New York, October 20-22, 2008.

David Heckerman. A Tutorial on Learning Bayesian Networks. *Technical Report 95-06*, Microsoft, 1995.

Ronald A. Howard. *Dynamic Programming and Markov Processes*. The M.I.T. Press, 1960.

Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge: Cambridge University Press, 1976.

Maria-Pina Limongelli, Eleni Chatzi, Michael Döhler and Geert Lombaert. Towards Extraction of Vibration-based Damage Indicators. *Factsheet of COST TU1402: Quantifying the Value of Structural Health Monitoring,* 2016

Chunlu Liu, Amin Hammad and Yoshito Itoh. Multi-objective optimization of bridge deck rehabilitation using a genetic algorithm. *Computer-Aided Civil and Infrastructure Engineering*, 12:431–43, 1997.

Min Liu and Dan M. Frangopol, Bridge Annual Maintenance Prioritization under Uncertainty by Multi-objective Combinatorial Optimization. *Computer-Aided Civil and Infrastructure Engineering*, 20:343–353, 2005. doi:10.1111/j.1467-8667.2005.00401.x

Jesus Luque and Daniel Straub. Risk-based optimization of inspection strategies in structural systems. *Structural Safety*, under review.

Samer Madanat and Moshe Ben-Akiva. Optimal inspection and repair policies for infrastructure facilities. *Transportation science,* 28(1):55-62, 1994.

Milad Memarzadeh and Matteo Pozzi. Integrated inspection scheduling and maintenance planning for infrastructure systems. *Computer-Aided Civil and Infrastructure Engineering*, 31(6):403-415, 2016.

Ayaho Miyamoto, Kei Kawamura and Hideaki Nakamura, H. Bridge management system and maintenance optimization for existing bridges. Computer-Aided Civil and Infrastructure Engineering, 15:45–55, 2000.

Andrew Y. Ng, A., and Michael Jordan. PEGASUS: A policy search method for large MDPs and POMDP. *In Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence (UAI)*, Stanford University, pages 406-415, Stanford, CA, June 30–July 3, 2000.

Konstantinos G. Papakonstantinou and Masanobu Shinozuka. Probabilistic model for steel corrosion in reinforced concrete structures of large dimensions considering crack effects, *Engineering Structures*, 57:306–26, 2013.

Konstantinos G. Papakonstantinou and Masanobu Shinozuka. Planning structural inspection and maintenance policies via dynamic programming and Markov processes. *Part I: Theory. Reliability Engineering & System Safety*, 130(0):202-213, 2014a.

Konstantinos G. Papakonstantinou and Masanobu Shinozuka. Planning structural inspection and maintenance policies via dynamic programming and Markov processes. *Part II: POMDP implementation. Reliability Engineering & System Safety*, 130(0):214-224, 2014b.

Konstantinos G. Papakonstantinou and Masanobu Shinozuka. Controlling hidden stochastic processes in real-time to optimize long-term infrastructure objectives. *Computational Stochastic Mechanics*. *Proceedings of the 7th International Conference (CSM-7)*, G. Deodatis and P.D. Spanos (eds.), Santorini, Greece, June 15-18, 2014c.

Konstantinos G. Papakonstantinou, Charalampos P. Andriotis and Masanobu Shinozuka. POMDP Solutions in Monitored Structures. *Proceedings of IFIP 2016*, 2016 working conference for the IFIP Working Group 7.5 on Reliability and Optimization of Structural Systems, Carnegie Mellon University, Pittsburgh, Pennsylvania, May 18-20, 2016.

Josep Porta, Matthijs Spaan and Nikos Vlassis. Value iteration for continuous state POMDPs. Technical report IAS-UVA-04-04, University of Amsterdam, Amsterdam, Netherlands, 2004.

Pascal Poupart and Craig Boutilier. VDCBPI: an approximate scalable algorithm for large scale POMDPs. *Advances in Neural Information Processing Systems*, 17:1081-1088, 2004.

Thomas L. Saaty. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York: McGraw-Hill, 1980.

Roland Schöbi and Eleni N. Chatzi. Maintenance planning using continuous-state partially observable Markov decision processes and non-linear action models. *Structure And Infrastructure Engineering*, 12(8):977-994, 2016.

Edward J. Sondik. *The optimal control of partial observable Markov processes*. Thesis (PhD). Stanford University, 1971.

John D. Sørensen. Framework for risk-based planning of operation and maintenance for offshore wind turbines. *Wind Energy*, 12:493–506, 2009. doi:10.1002/we.344

Daniel Straub and Eleni Chatzi. Framework for SHM Value of Information Analysis and Computation. *Factsheet of COST TU1402: Quantifying the Value of Structural Health Monitoring*, 2016.

Daniel Straub, E. Chatzi, E. Bismut, W.M.G. Courage, M. Döhler, M.H. Faber, J. Köhler, G. Lombaert, P. Omenzetter, M. Pozzi, S. Thöns, D. Val, H. Wenzel, D. Zonta. Value of Information: A roadmap to quantifying the benefit of structural health monitoring, *In proceedings: ICOSSAR 2017 - 12th International Conference on Structural Safety & Reliability*, Vienna, Austria, 6-10 August 2017.