Reliability Engineering and System Safety 153 (2016) 159–169

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress

Hierarchical modeling of systems with similar components: A framework for adaptive monitoring and control

Milad Memarzadeh [a,*,1], Matteo Pozzi [a,2], J. Zico Kolter [b,3]

[a] Department of Civil and Environmental Engineering, Carnegie Mellon University, United States
[b] School of Computer Science, Carnegie Mellon University, United States

ARTICLE INFO

Article history:
Received 10 July 2015
Received in revised form
12 February 2016
Accepted 28 April 2016
Available online 7 May 2016

Keywords:
Adaptive monitoring and control
Markov Decision Process
Hierarchical Bayesian modeling
Markov Chain Monte Carlo

ABSTRACT

System management includes the selection of maintenance actions depending on the available observations: when a system is made up by components known to be similar, data collected on one is also relevant for the management of others. This is typically the case of wind farms, which are made up by similar turbines. Optimal management of wind farms is an important task due to high cost of turbines' operation and maintenance: in this context, we recently proposed a method for planning and learning at system-level, called PLUS, built upon the Partially Observable Markov Decision Process (POMDP) framework, which treats transition and emission probabilities as random variables, and is therefore suitable for including model uncertainty. PLUS models the components as independent or identical. In this paper, we extend that formulation, allowing for a weaker similarity among components. The proposed approach, called Multiple Uncertain POMDP (MU-POMDP), models the components as POMDPs, and assumes the corresponding parameters as dependent random variables. Through this framework, we can calibrate specific degradation and emission models for each component while, at the same time, process observations at system-level. We compare the performance of the proposed MU-POMDP with PLUS, and discuss its potential and computational complexity.

1. Introduction and previous works

Infrastructure systems age and their condition deteriorate due to fatigue and environmental loads. Accurate risk analysis is crucial to extend their life-span, and to guide decision making towards a sustainable use of resources. This analysis should be conducted in a probabilistic framework, modeling the degradation process and incorporating the effect of the maintenance actions. Wind farms are one of the many type of infrastructures that is recently growing due to the need for sustainable renewable energy. On average, operation and maintenance (O&M) costs account for about 25–30% of the overall expenses for wind energy generation [1]; however, a careful optimization process can reduce those costs and make wind sector more competitive with other renewable sources. Specifically, the wind energy infrastructure is exposed to fast deterioration, and requires frequent maintenance. However, a large amount of data related to system deterioration and performance is available, and allows for an improvement of the O&M process. An accurate, efficient and reliable framework to optimally manage wind farms includes selecting appropriate actions to balance between exploration of the degradation behavior and a profitable exploitation, integrating information on components (i.e. the turbines) at the level of the system (i.e. the farm).

Methods based on Markov Decision Processes (MDP) have been used for O&M of infrastructures systems both at component and system level [2–6]. A Partially Observable Markov Decision Process (POMDP) is a generalization of the MDP [7,8], where the exact state of a component cannot be observed directly, but can only be inferred by indirect observations. Time is discretized in steps and the goal is, at each step, to select a maintenance action to maximize the "value" i.e. the sum of discounted expected rewards over the infinite horizon. The POMDP framework has been used in recent years for optimal management of wind farms [9–11]. Specifically, Byon et al. [9] have developed a method accounting for uncertain weather conditions in the maintenance strategy, and they have extended their original method to season-dependent condition-based maintenance [10]. Nielsen and Sorensen [11] have investigated the use of limited-information influence diagrams to support decision making for O&M of offshore turbines. McMillan and Ault [12] have evaluated the effect of adopting MDP in

* Corresponding author.
E-mail addresses: miladm@cmu.edu (M. Memarzadeh), mpozzi@cmu.edu (M. Pozzi), zkolter@cs.cmu.edu (J.Z. Kolter).
[1] Tel.: +1 540 557 7087.
[2] Tel.: +1 412 268 5649.
[3] Tel.: +1 412 268 4733.

http://dx.doi.org/10.1016/j.ress.2016.04.016
0951-8320/

modeling the deterioration of turbines, using Monte Carlo simulations to assess the cost effectiveness of condition-based monitoring.

A limit in this scientific literature is that the transition probabilities (modeling the degradation behavior of the turbines) and emission probabilities (modeling the precision of the monitoring system) are assumed as fixed values, and epistemic uncertainty is not taken into account. However, in most real-world management of infrastructure systems, high uncertainty may affect these parameters.

Recent works such as that of Ross et al. [13] have generalized the POMDP framework to Bayes-Adaptive POMDP (BA-POMDP) that allows treatment of the model parameters as random variables, whose distribution can be learned and updated during the decision process. However, the computational complexity of this framework is high and it becomes easily intractable with increase in the time horizon or dimensionality (i.e. number of states, observations, and actions). Jaulmes et al. [14,15] have proposed an algorithm called "MEDUSA" (Markovian Exploration with Decision based on the Use of Sampled models Algorithm) to approximately identify the optimal policy when the model is unknown: their algorithm updates the model incrementally using selected queries, while still optimizing the estimate value. Recently, we have proposed an alternative method for approximate learning and planning, named "PLUS" (Planning and Learning for Uncertain Dynamic Systems) [16,17], which performs better than the previous method, in our setting. Specifically, that algorithm can process observations, update the distributions of model parameters, and select the optimal strategy accounting for model uncertainty. PLUS is structured in two phases: learning and planning. In the learning phase, it uses Markov Chain Monte Carlo (MCMC) Gibbs sampling [18], while the planning phase relies on an approximation which neglects the exploratory value of learning the model parameters: details about PLUS algorithm can be found in Memarzadeh et al. [17]. The approach allows for a rational treatment of collected data (e.g. by sensors and visual inspections), a probabilistic tracking of the components' condition, and robust decision making support. There are two modes of implementing PLUS at system-level: the first one, which we call *Individual PLUS*, assumes that components are completely independent from each other. The second, that we name *Global PLUS*, assumes that components are modeled by identical stochastic processes.

Still, PLUS cannot model the intermediate case of components that are known to be similar but not identical. In this paper, we extend that method by proposing a hierarchical Bayesian modeling approach that, at a higher computation cost, is able to model a system made up by similar components, and allows a system-level flow of information without forcing the stochastic processes to be identical. Section 2 describes the problem statement; Section 3 introduces the general MU-POMDP framework and the proposed method for learning model parameters and hyper-parameters; Section 4 investigates the performance of the method on a low-dimensional problem, Section 5 applies it to a wind farm management problem, and Section 6 draws conclusions.

2. Problem statement

Our motivating application is O&M of a wind farm made by similar turbines. In that context, the degradation behavior and the accuracy of the monitoring system are uncertain, but can be assumed to be similar across components, e.g. because components of different typologies are exposed to the same environment, or components of the same typology are exposed to different environments.

Formally, we define the problem as follow. Suppose to manage a set of K components, each modeled by a POMDP. The set of parameters controlling the POMDPs are uncertain and dependent. In this context, how can we (i) formulate a probabilistic model to capture the dependence among the components, (ii) develop an analytical and numerical technique to infer the variables in the problem, and (iii) define an approach to identify the optimal management policy? While the limit cases of independent and of identical models can be solved by Individual and Global PLUS respectively, the intermediate setting poses specific computational problems that we address in this paper.

3. Proposed methodology

3.1. MU-POMDP framework

To address the first question posed in previous section, we make use of the hierarchical Bayesian modeling approach. We refer to the proposed framework as Multiple Uncertain POMDP (MU-POMDP), and Fig. 1 shows the corresponding probabilistic graphical model, or "influence diagram", for a system with two components ($K=2$). Only variables related to time steps $(t-1)$ and (t) are shown in the figure. The reader is referred to Kaelbling et al. [20] and to [17] for technical details of the classical POMDP framework which, as indicated in the figure, is used to model each component. The graphical model follows the notation of dynamic Bayesian network and influence diagrams adopted from the textbook of Barber [21]: circles define random variables, squares decision variables, diamonds utility variables, dots fixed parameters, and arrows define dependence among variables; shaded variables are observed, and a dashed line indicates that an observation is available before taking an action. S indicates the component state, A the maintenance action, Z the available observation and R the monetary loss. Subscript "k,t" refers the variable to component k at time t. We assume all components have the same number of possible states, actions and observations, and we indicate these quantities as $|S|$, $|A|$ and $|Z|$ respectively. At any time, the manager collects current observations and select actions for the next time step. $\mathbf{T}_k$ indicates the transition probability of component k, which is a 3-dimensional matrix defined over the space of current states, actions, and future states $[|S|\times|A|\times|S|]$ and $\mathbf{O}_k$ indicates the corresponding emission probability, which is a 3-dimensional matrix defined over the space of states, actions, and observations $[|S|\times|A|\times|Z|]$. Transition and emission probabilities are the "model parameters" and are formally defined by conditional probabilities, as $T_k(S_{k,t-1},A_{k,t-1},S_{k,t})=P(S_{k,t}|S_{k,t-1},A_{k,t-1},\mathbf{T}_k)$ and $O_k(S_{k,t},A_{k,t-1},Z_{k,t})=P(Z_{k,t}\ S_{k,t},A_{k,t-1},\mathbf{O}_k)$ respectively. They are fixed parameters in the POMDP framework, and treated as random variables in the BA-POMDP one.

MU-POMDP includes an additional layer of hyper-parameters, to capture the dependence among the model parameters of different components. Hyper-parameters are marked as α_T, $\boldsymbol{\beta}_T$, α_O and $\boldsymbol{\beta}_O$ in Fig. 1: the first two values define the dependence in the transitions, while the latter define that of emissions. While model parameters are different for each component, hyper-parameters are common to the entire system. Formally, matrices $\boldsymbol{\beta}_T$ and $\boldsymbol{\beta}_O$ have the same dimension of $\mathbf{T}_k$ and $\mathbf{O}_k$ respectively, while α_T and α_O are scalar variables. The role of these variables will become apparent in the following sections. Parameter matrices $\boldsymbol{\eta}_T$ and $\boldsymbol{\eta}_O$, of dimension equal to that of $\mathbf{T}_k$ and $\mathbf{O}_k$ respectively, and scalar variables λ_T and λ_O define the distribution of hyper-parameters.

3.2. Probabilistic inference via MCMC

The overall purpose of the inference task is to represent the posterior distribution of the variables in the process. In this context, the posterior distribution is defined conditional to all observations Z and actions A observed up to present time. We use the upper bar to