Reliability Engineering and System Safety 172 (2018) 45–57

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

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

Value-of-information in spatio-temporal systems: Sensor placement and scheduling

C. Malings*, M. Pozzi

Department of Civil & Environmental Engineering, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA

A B S T R A C T

The management of infrastructure involves accounting for factors which vary in space over the system domain and in time as the system changes. Effective system management should be guided by models which account for uncertainty in these influencing factors as well as for information gathered to reduce this uncertainty. In this paper, we address the problem of optimal information collection for spatially distributed dynamic infrastructure systems. Based on prior information, a monitoring scheme can be designed, including placement and scheduling of sensors. This scheme can be adapted during the management process, as more information becomes available. Optimality can be defined in terms of the value of information (VoI), which provides a rational metric for quantifying the benefits of data gathering efforts to support system management decision-making. However, the computation of this metric in spatially and temporally extensive systems can present a practical impediment to its implementation. We describe this complexity, and investigate a special case of system topology, termed as a temporally decomposable system with uncontrolled evolution, in which the complexity of assessing VoI grows at a manageable rate with respect to the system management time duration. We demonstrate the evaluation and optimization of the VoI in an example of such a system.

1. Motivation and background

In this paper, we examine the optimization of sensor placements and scheduling to support the management of infrastructure systems. For the management of large systems with numerous components whose states evolve in time, the determination of optimal policies both for maintaining these systems and for inspecting these systems to determine what management actions are needed are important questions [1–5]. As a motivating example of sensing and decision-making in an evolving system, consider a system whose performance is influenced by a physical quantity which varies in both space and time, such as depicted in Fig. 1; a square region, described by horizontal coordinates x_1 and x_2, is represented at different time instants $t = 1$ and $t = 2$. The set of random values for this physical quantity at each coordinate in space and time is described as the random field. For instance, this field might represent the temperature to which a population is subjected, which can cause health difficulties, placing strain on a medical system e.g., [6,7]. An underlying probabilistic model of the random field captures its characteristics, including its expected value, variability, and interdependence relationships in space and time e.g. [8]. In Fig. 1a, a map of the field at one time is depicted, showing its variation in space. Fig. 1c depicts the random field at a subsequent time; note the similarities in the field shape, stemming from the modeled correlation of the field across time. The infrastructure system might then consist of a continuous domain, e.g. in the case of population in a region exposed to extreme temperatures, or of discrete components occupying positions within this domain, e.g. in the case of several critical assets vulnerable to high temperature.

Based on prior knowledge, where the field is predicted to exceed a threshold, an appropriate intervention activity can be carried out to avoid the negative consequences of this exceedance. For example, heat advisories might be issued for specific regions and times, which can mitigate the consequences of the population being exposed to extreme heat. Intervention decisions take into account the consequences of different possible outcomes, the costs of response options, and the inherent uncertainty in the field. Measurements of the field can also be made and used to update the prior model. For example, in Fig. 1a, x's indicate locations where measurements of random field values are made, and Fig. 1b presents an example of decision-making based on these measurements. The red area denotes where the field exceeds a set threshold (causing a local failure in the system), while the blue area depicts where, based on the updated knowledge of the field obtained by processing the available measures, threshold exceedance is predicted, and therefore an appropriate intervention is taken. Note that there is not perfect overlapping between these regions, since there remains residual uncertainty in the field.

Information can be costly to acquire, and therefore should be prioritized in both time and space to trade off the costs of collecting this information against its potential benefits. In Fig. 1a, measurements at initial time $t = 1$ are distributed evenly over the domain of the system, to provide adequate spatial coverage. In Fig. 1c, at a later time, measures are again distributed evenly, but at different locations; this reflects the temporal correlation of the field, which makes repeated measures at the

* Corresponding author.
E-mail address: cmalings@andrew.cmu.edu (C. Malings).

https://doi.org/10.1016/j.ress.2017.11.019
Received 30 May 2017; Received in revised form 15 October 2017; Accepted 24 November 2017
0951-8320/

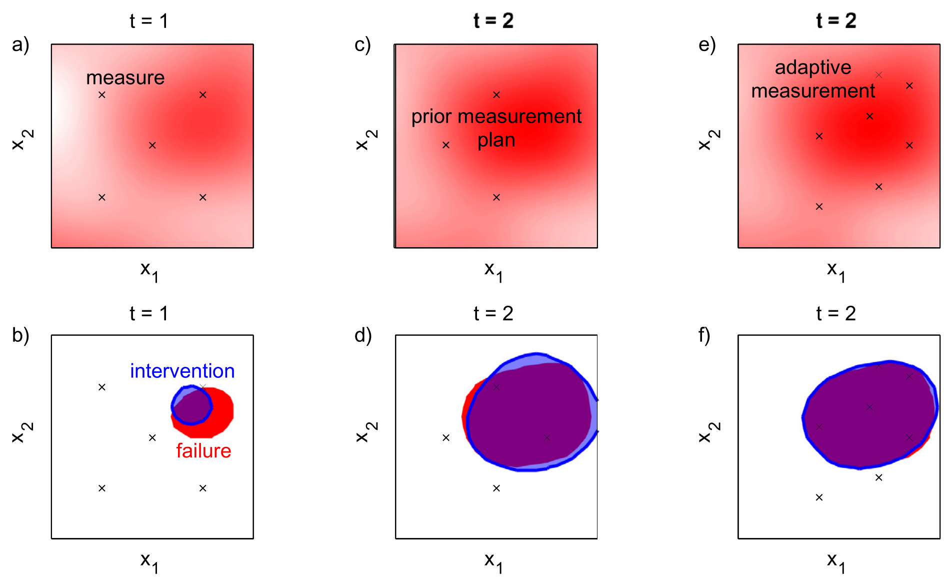

Fig. 1. A motivating example of the monitoring and management of a system in space and time.

same location somewhat redundant. This represents the related problems of sensor placement, i.e., determining an appropriate spatial arrangement of sensors, and of scheduling, i.e., determining at what times measures should be collected.

Decisions about sensor placement and scheduling take a variety of factors into account. High prior uncertainty in the random field can be reduced through sensing. However, factors relevant to decision-making, such as the likelihood and potential consequences of making an incorrect decision without additional data, should also be accounted for. Furthermore, the interdependence structure of the random field, in both space and time, should inform the sensing plan. In space, collecting many measures in highly interdependent fields can be redundant, but in fields with weak interdependence, more closely spaced measures may be necessary to avoid missing features of interest. In time, earlier measurements can help to identify trends and support later decision-making, but need also be updated as information becomes out-of-date. Finally, the relative precisions and costs of different sensors, especially costs relating to sensor placement (whether it is cheap or expensive to gather a measurement at a new location) and scheduling (whether it is cheap or expensive to repeatedly collect measurements at the same location) should be taken into account when determining which measures will be cost-effective.

Finally, there is the problem of online or adaptive sensing, where sensor placements and schedules can be changed in light of new information. Fig. 1e depicts such a case, where, because of the high observed random field values in the upper right, at a later time more measurements are allocated for this area to better determine whether or not the field will exceed the threshold. By comparing Fig. 1d and Fig. 1f, the greater number and concentration of measurements in the upper-left allows the intervention zone to more closely match the area of exceedance. This illustrates the potential benefits of adaptive or online sensing, but these benefits should be traded off against the additional costs of evaluating and implementing a revised sensing plan.

In this paper, we examine how to optimally place and schedule measurements to best support decision-making for system management by taking into account the various factors mentioned above. We do this making use of the value of information (VoI) to explicitly trade off the benefits of collected information, in terms of improved decision-making, against the costs of information collection [9]. The VoI metric aligns well with many of the intuitive ideas discussed above of what makes a sensor placement in space informative to system management [10]. Here, we extend these results from static sensor placements to sensor placements and schedules in dynamic systems, and also to adaptive sensing, where sensor placements can change over time. Previous work has made use of VoI to quantify the benefits of structural health monitoring efforts [11–13], optimize the positioning of sensors to support the management of structures under uncertain extreme loading [10,14], and optimize inspection schemes for deteriorating components [15,16]. Other approaches to sensor scheduling, making use of concepts such as observability and state estimation accuracy, have also been applied to this problem e.g., [17,18].

In general, the computational cost of VoI evaluation grows exponentially as the size of the system increases; this can be seen by examining the management decision-making problem via a decision tree and noting that the number of “leaves” will grow exponentially as the number of possible management actions, observations, and system states increases [19]. Previous work in spatial systems identified a special case of system topology, termed as a cumulative system topology, in which this exponential growth can be reduced to a linear growth in the number of system components, provided management activities are conducted locally for each component [14].

In Section 2 of this paper, the assumption that the actions taken to manage an evolving system have temporally local effects is used to identify a corresponding special case in which the computational demand of VoI evaluation is merely linear in the time duration for system management. In Section 3, we give a brief overview of greedy offline and online approaches to efficient sensor placement and scheduling based on VoI which are used in this paper; further information on these methods is available in the companion paper [20]. In Section 4, we introduce a Gaussian random field modeling framework for spatio-temporal systems. This framework is used in Section 5 to demonstrate the application of the VoI metric in two examples: a simulated system modeling differential settlement between structural columns over time and a problem based on measurements taken on structural columns during the construction of the Scott Hall building at Carnegie Mellon University. Finally, some general conclusions are drawn in Section 6.

2. Value of information in spatio-temporal systems

This section begins by outlining a model for the monitoring and management of a system whose behavior is affected by random variables which vary in both time and space in Section 2.1. Within this model, the VoI metric is defined in general in Section 2.2. The metric is also examined under several assumptions on the structure of the system and its management which lead to increasing computational tractability for the

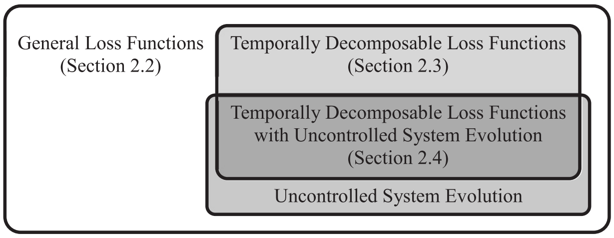

Fig. 2. Relationships of the assumptions presented in the following sections, which present special sub-cases of the general loss function definition.

metric. The relationships of these assumptions are depicted in Fig. 2. The assumption that the loss function is temporally decomposable allows for more efficient evaluation of the VoI, as described in Section 2.3. An additional assumption states that the evolution of the system is uncontrolled, i.e., the evolution of the underlying random fields affecting the system is not influenced by actions taken for system management. Where this assumption overlaps with the previous, as discussed in Section 2.4, the computational demand of evaluating the VoI is reduced to be linear in the length of the time duration for system management. This represents an interesting special case of system management (of which the example discussed in Section 1 is a representative) in which VoI can be used in a computationally tractable way to support sensor placement and scheduling.

2.1. Probabilistic spatio-temporal system model

Let F indicate a random field which affects the performance of an infrastructure system over management time duration T. Let T be discretized into m timesteps, denoted $T = \{t_1, \dots, t_m\}$. Let $\mathrm{f}(\mathbf{x}_i, t_j)$ denote the ith of n random variables which affect the system at timestep t_j, where the coordinate $\mathbf{x}_i$ is used as an index or reference coordinate for the variable, e.g. if the variable is associated with a particular factor at a particular spatial location, then its coordinate will correspond with this spatial location. In the case of multiple co-located random field affecting a system (e.g. a temperature and humidity field affecting corrosion), this coordinate is augmented with an indexing term to distinguish between co-located fields. The spatio-temporal random field which affects the system should be described by an appropriate spatio-temporal random field model which captures the prior knowledge of the distributions of random variables, including their uncertainties, spatial interdependencies, and temporal evolution via proper spatio-temporal joint probability functions. The selection of such functions represents an important problem in itself. Where possible, this selection can be performed by conducting numerical simulations using a deterministic physics-based model of the system, using empirical data collected in similar systems to the one being modeled, or soliciting expert judgments on the most appropriate model forms. The reader is referred to [8] for a comprehensive overview of approaches to the spatio-temporal model selection problem, and to [6] for a recent application of these approaches to define a probabilistic model of an environmental hazard impacting an urban system. Alternatively, hierarchical modeling can be used in which the functional forms and parameters of the random field are described by probability distributions. This allows for additional flexibility in the model but greatly increases the computational cost associated with its use [21].

Let $\mathbf{f}$ denote the vector of random variables affecting the system at all discrete timesteps over time duration T. This vector can be expressed as $\mathbf{f} = [\mathbf{f}_1{}^{\mathrm{T}}, \dots, \mathbf{f}_m{}^{\mathrm{T}}]^{\mathrm{T}}$, with sub-vector $\mathbf{f}_j$ denoting the random variables acting on the system at time t_j. The prior distribution for the vector of random variables $\mathbf{f}$ is denoted as p_F, with p_{F_j} indicating the prior distribution for $\mathbf{f}_j$.

Let Y denote a plan to measure the variables affecting the system over the management time duration, i.e. a sensor placement and scheduling scheme for the system. Vector $\mathbf{y}$ denotes a specific outcome of this scheme, i.e. a set of possibly noisy measures collected on the system, with $\mathbf{y}_j$ indicating a subset of measures which are first available to the decision-making agent at time t_j. Thus, indexing of the measures indicates not necessarily when they are taken or which time index the measured random variables are associated with, but when the measures first become available to support decision-making; for example, measurement set $\mathbf{y}_j$ might correspond to observations of variables $\mathbf{f}_{j-1}$ if there is a one timestep delay between when measures are collected and when they can be processed to support decision-making.

Based on these measurements, the prior distribution of random variables can be updated to a posterior distribution. This updating can be performed using standard techniques for Bayesian inference. This posterior distribution after observation of $\mathbf{y}$ is denoted as $\mathrm{p}_{F|\mathbf{y}}$. Note, however, that inference and updating can be performed using partial information, i.e., at time t_j, the prior distribution can be updated using all information available up to this time. We denote by $Y_{\to j}$ the subset of measurements whose outcomes will be available to the managing agent at time t_j, with $\mathbf{y}_{\to j}$ denoting the specific observations obtained up to and including that time. Thus, at time t_j, the most up-to-date posterior distribution for the random variables affecting the system, utilizing all information collected on the system which is available up to that point, is $\mathrm{p}_{F|\mathbf{y}_{\to j}}$.

In managing an uncertain system, decision-makers select actions to take to intervene in the system. Let $\mathbf{a}$ denote a set of selected actions for managing the system over the entire time duration, selected from set $\mathcal{A}$ of all possible action sequences (including a null or 'do nothing' sequence). Let $\mathbf{a}_j$ denote the subset of actions decided upon at the jth timestep of the management duration. Note that this timestep need not necessarily denote when the actions are implemented, only the last timestep in which the actions are free to be altered, i.e. the timestep at which the choice of actions is 'locked in' and cannot be changed later.

Actions taken to manage the system can have an impact on the evolution of the random field underlying the system. Thus, the distribution of the random variables is in general denoted as $\mathrm{p}_{F|\mathbf{a}}$. We denote the prior distribution p_F as a special case where a sequence of actions designated as the null sequence is taken, i.e. no interventions are taken to affect the evolution of the system. To preserve causality, it is assumed that actions decided upon at the jth timestep can affect the evolution of the system from timestep $j+1$ onward, i.e. it takes a minimum of one discrete timestep for the impacts of actions to have an effect of the random variables which describe the future performance of the system. Also note that the effects of future actions cannot propagate backward in time, affecting the state of the system in earlier timesteps. Formally, we have that $\mathrm{p}_{F_{\to j}|\mathbf{a}_{j\to}} = \mathrm{p}_{F_{\to j}} \, \forall \, \mathbf{a}_{j\to} \in \mathcal{A}_{j\to}$, i.e. that the distribution for random variables affecting the system up to and including time t_j is the same regardless of the choice of actions taken from that timestep onward (denoted $\mathbf{a}_{j\to}$).

The last element needed to define the system management decision-making problem is the loss function, denoted $\mathrm{L}(\mathbf{f}, \mathbf{a})$, which represents a mapping from the variables affecting the system and the actions taken to manage it to a scalar quantity representing the utility of that outcome to the system's managing agent. Typically, this is expressed in monetary terms as the lifetime cost (or negative revenue) of managing the system over the time duration T.

The loss function captures all costs which are relevant to the decision-making problem resulting from certain combinations of variable states and actions. These include the cost of taking management actions, such as repairing a potentially damaged component or closing down part of a system for safety reasons. They also include the cost of potential failures or reductions in system performance, such as the cost of lost revenues due to system down-time, costs or penalties of failing to meet serviceability requirements, and costs of property damage and potential loss-of-life in the case of catastrophic failure. The loss function also captures instantaneous effects of management actions on the system; while it is assumed that actions $\mathbf{a}_j$ cannot affect variables $\mathbf{f}_j$, different combinations of actions and variables at the same timestep can have different consequences, as captured by the loss function. We assume that the loss function does not capture the costs of monitoring the system; these are taken into account by a separate sensing cost function C(Y), so that sensing and management costs can be examined separately and traded off. In Sections 2.2 through 2.4, various assumptions on the structure of the loss function and the effects of management actions on the system are used to define the VoI of observations in various contexts.

2.2. General loss functions

First, we consider the case of a general loss function with no special properties. The loss incurred by the system manager is, under this most general definition, a function of all variables **f** affecting the system and all actions **a** taken for system management.

Without access to any measurements of the system, the best an infrastructure manager can hope to do is select a set of management actions which will minimize (in an expected sense) the loss function. This is termed the prior expected loss, and is evaluated as:

$$\mathbb{E}\mathrm{L}(\emptyset) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{a}}\mathrm{L}(\mathbf{f},\mathbf{a}) \tag{1}$$

where $\mathbb{E}_{F|\mathbf{a}}$ denotes the statistical expectation with respect to $\mathrm{p}_{F|\mathbf{a}}$. There is no advantage to waiting to select management actions, as no additional information will be available to alter the decision-making process, and so the entire set of optimal actions for managing the system throughout the time duration can be selected at once; this set (also referred to as the "open-loop policy") is denoted $\mathbf{a}^*(\emptyset)$, and is the argument which minimizes Eq. (1).

If, on the other hand, measurements of the system are available to the managing agent during the time duration, choices of actions can and should be changed based on new information. The standard approach to this type of sequential information collection and decision making problem is dynamic programming [22]. We follow this approach to define the posterior expected loss $\mathbb{E}\mathrm{L}(Y)$, or the expected loss for managing the system with access to measurements from the set Y. We define this recursively, beginning with the final timestep of system management and working backward. In timestep m, the complete set of measurements $\mathbf{y}$ is available, and the record of past (or implemented) management actions $\mathbf{a}_{\rightarrow(m-1)}$ is known. Based on this information, the agent should select a set of actions for managing the system at the final timestep, $\mathbf{a}_m^*$, which minimizes the expected loss given all observations, all past actions, and the current action:

$$\mathbf{a}_m^*\left(\mathbf{a}_{\rightarrow(m-1)},\mathbf{y}\right) = \mathrm{argmin}_{\mathbf{a}_m\in\mathcal{A}_m} \mathbb{E}_{F|\mathbf{a},\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \tag{2}$$

where $\mathbf{a} = \{\mathbf{a}_{\rightarrow(m-1)}, \mathbf{a}_m\}$ concatenates the past and present actions.

We now introduce a value function which defines the expected cost to manage the system given all implemented actions and collected observations before that timestep. For timestep m, the value function is:

$$\mathbb{E}\mathrm{L}_m^*\left(\mathbf{a}_{\rightarrow(m-1)},\mathbf{y}_{\rightarrow(m-1)}\right) = \mathbb{E}_{Y_m|\mathbf{a}_{\rightarrow(m-1)},\mathbf{y}_{\rightarrow(m-1)}} \min_{\mathbf{a}_m\in\mathcal{A}_m} \mathbb{E}_{F|\mathbf{a},\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \tag{3}$$

That is, it is the expectation (over the final set of measures) of the minimum (over the final set of actions) of the expected loss for managing the system. Expectations are conditioned on all available information and all previous action choices.

The value function can be defined recursively by noting that, at each timestep, a set of actions are selected to minimize the expected loss, and a new set of observations are taken (which are influenced by all past information and actions). Therefore, the recursive definition of the value function is:

$$\mathbb{E}\mathrm{L}_j^*\left(\mathbf{a}_{\rightarrow(j-1)},\mathbf{y}_{\rightarrow(j-1)}\right) = \mathbb{E}_{Y_j|\mathbf{a}_{\rightarrow(j-1)},\mathbf{y}_{\rightarrow(j-1)}} \min_{\mathbf{a}_j\in\mathcal{A}_j} \mathbb{E}\mathrm{L}_{(j+1)}^*\left(\mathbf{a}_{\rightarrow j},\mathbf{y}_{\rightarrow j}\right) \tag{4}$$

while in the final timestep the value function has a special, explicit definition:

$$\mathbb{E}\mathrm{L}_{(m+1)}^*\left(\mathbf{a}_{\rightarrow m},\mathbf{y}_{\rightarrow m}\right) = \mathbb{E}_{F|\mathbf{a}_{\rightarrow m},\mathbf{y}_{\rightarrow m}}\mathrm{L}\left(\mathbf{f},\mathbf{a}_{\rightarrow m}\right) \tag{5}$$

The posterior expected loss under measurement scheme Y is the value function at the first timestep:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}\mathrm{L}_1^*\left(\mathbf{a}_{\rightarrow 0},\mathbf{y}_{\rightarrow 0}\right) \tag{6}$$

where $\mathbf{a}_{\rightarrow 0}$ and $\mathbf{y}_{\rightarrow 0}$ are by definition empty. Using the recursive definition of Eq. (4), Eq. (6) can be expanded as:

$$\mathbb{E}\mathrm{L}(Y) = \mathbb{E}_{Y_1}\min_{\mathbf{a}_1\in\mathcal{A}_1}\mathbb{E}_{Y_2|\mathbf{a}_1,\mathbf{y}_1}\min_{\mathbf{a}_2\in\mathcal{A}_2}\cdots\mathbb{E}_{Y_m|\mathbf{a}_{\rightarrow(m-1)},\mathbf{y}_{\rightarrow(m-1)}} \min_{\mathbf{a}_m\in\mathcal{A}_m}\mathbb{E}_{F|\mathbf{a},\mathbf{y}}\mathrm{L}(\mathbf{f},\mathbf{a}) \tag{7}$$

Here it is clear that the posterior expected loss involves repeated nesting of expectations (with respect to observations) and minimizations (with respect to actions) down the decision tree. This parallels the sequential collection of information and determination of actions which is necessary for system management.

In general, in timestep j, actions should be selected so as to minimize the value function for the next timestep, a function of all information collected and all actions taken up to that point:

$$\mathbf{a}_j^*\left(\mathbf{a}_{\rightarrow(j-1)},\mathbf{y}_{\rightarrow j}\right) = \mathrm{argmin}_{\mathbf{a}_j\in\mathcal{A}_j}\mathbb{E}\mathrm{L}_{(j+1)}^*\left(\mathbf{a}_{\rightarrow j},\mathbf{y}_{\rightarrow j}\right) \tag{8}$$

For a measurement outcome $\mathbf{y}$, beginning from the first timestep and moving to the last, a sequence of optimal actions $\mathbf{a}^*(\mathbf{y})$ is defined via iterative solution of Eq. (8).

The VoI of measurement scheme Y is defined as the difference between the prior and posterior expected losses under this scheme, i.e. as the expected reduction in loss due to the taking of more appropriate actions based on the gathered information [9]:

$$\mathrm{VoI}(Y) = \mathbb{E}\mathrm{L}(\emptyset) - \mathbb{E}\mathrm{L}(Y) \tag{9}$$

Evaluation of the VoI for a given observation scheme can be a computationally daunting task in spatio-temporal systems. It involves, for all possible measurement outcomes $\mathbf{y}$, determining an optimal action sequence $\mathbf{a}^*(\mathbf{y})$ and evaluating the expected loss under this action sequence. We quantify the problem dimensionality in a reference case where random variables are discrete, having θ_f possible outcomes, there are θ_a possible choices for each management action in $\mathcal{A}$, and each observation in Y has θ_y possible results. The computational bottleneck is the evaluation of Eq. (5) for every sequence of observations and actions. Therefore, using order notation, the search space of general VoI evaluation is $\mathcal{O}(\theta_y^{|Y|}\theta_a^{|\mathcal{A}|}\theta_f^{mn})$, representing the fact that, for each of $\theta_y^{|Y|}$ possible observation sequences and $\theta_a^{|\mathcal{A}|}$ possible management action sequences, the loss must be averaged across θ_f^{mn} possible random field variable states to determine the value function. If we further assume that variables, actions, and observations are binary, i.e. that $\theta_y = \theta_a = \theta_f = 2$, and that at each timestep n_y observations are made and n_a actions are taken, i.e. the number of observations and actions is linear with respect to the management time duration, then $|Y| = mn_y$ and $|\mathcal{A}| = mn_a$, and the dimensionality can be expressed as $\mathcal{O}(2^{m(n_y+n_a+n)})$.

The computational complexity of VoI assessment under a general loss function via the methods outlined above is intractable for large systems [23,24]. Although the above example of complexity growth focused on discrete states, observations, and actions, the complexity growth rate is analogous in problems with continuous states, observations, and actions. Note also that for systems with continuous action variables, while efficient solution to the minimization problem over actions might be possible where the problem structure is convex, there is in general no guarantee that this will be the case, and computationally intensive non-convex optimization techniques would need to be employed. Many approximate solution approaches have been developed to address this problem of complexity growth. Sampling-based approaches, including Monte Carlo and Markov Chain Monte Carlo techniques, have been used for performing Bayesian inference and approximating VoI. However, appropriate selection of sample size is important to avoid biased estimates of the VoI, and many samples may be required to produce a suitably accurate result in large systems [25]. Surrogate or emulator models have also

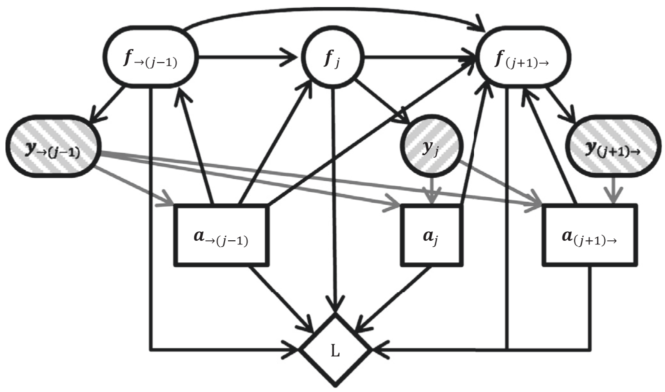

Fig. 3. Probabilistic graphical model for a system with a general loss function. Potential interdependencies between measurements have been removed to improve clarity.

been used, approximating VoI in large problems using more easily evaluated functions [26,27]. Hierarchical dynamic Bayesian models with associated approximate inference methods have also been used [28–31]. In this paper, we develop an approach to tractable and exact evaluation of VoI in large evolving systems by exploiting particular structures which the loss function might have. These loss function structures, their impacts on VoI evaluation complexity, and the consequences of these structures in terms of assumptions about the type of system being modeled, are introduced and discussed in Sections 2.3 and 2.4.

Fig. 3 represents a probabilistic graphical model (PGM) of a system with a general loss function. This follows a common convention for PGMs, with circles representing random variables, shaded circles (or ovals) representing observed variables, squares (or rectangles) representing inputs or decisions, rhombi representing deterministic outcomes, and lines and arrows indicating probabilistic or deterministic relationships among variables [19]. Random variables $\mathbf{f}$, observations $\mathbf{y}$, and actions $\mathbf{a}$ are split into three categories: those before timestep j (i.e. the past), those at this timestep (i.e. the present), and those after this timestep (i.e. the future). Note that past actions and variable states affect present and future variable states, and that present variable states and actions affect future variable states. Also note that present actions are made with knowledge of past and (depending on the problem context) present observations, as indicated by the shaded arrows. The loss, in general, is a function of past, present, and future actions and variable states.

2.3. Temporally decomposable loss functions

We now assume that the loss function is decomposable across time as follows:

$$\mathrm{L}(\mathbf{f},\mathbf{a}) = \sum_{j=1}^{m} \gamma_j \mathrm{L}_j\left(\mathbf{f}_j,\mathbf{a}_j\right) \tag{10}$$

That is, the total loss is expressed as the discounted sum of losses associated with each timestep, where loss $\mathrm{L}_j(\mathbf{f}_j, \mathbf{a}_j)$ is associated with the jth timestep and γ_j is the positive discounting factor associated with this timestep. Such loss functions are commonly encountered in engineering applications where the costs of system failures and of executing actions are associated with specific timesteps and are discounted back to their present value to evaluate the lifetime loss for the system. This decomposable form for the loss function is also commonly used for partially observable Markov decision processes, or POMDPs, e.g. [16]. Note, however, that in this case the Markovian assumption of POMDPs is not applied, and that a more general discounting scheme is used. Fig. 4 depicts a PGM for a system with a decomposable loss function.

The assumption of a decomposable loss function, together with the linearity of the expectation, allows for the evaluation of expected loss to be performed in a different and more efficient way. For example,

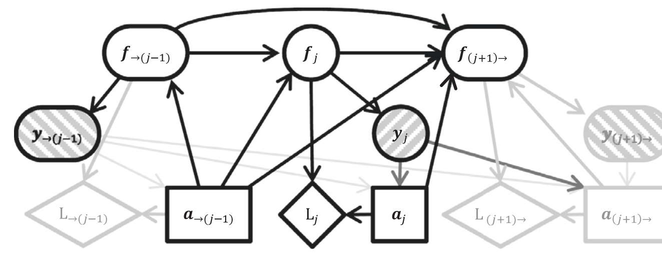

Fig. 4. Probabilistic graphical model for a system with a decomposable loss function. Potential interdependencies between measurements have been removed to improve clarity.

Eq. (5) can now be expressed as:

$$\mathbb{E}\mathrm{L}^*_{(m+1)}\left(\mathbf{a}_{\to m},\mathbf{y}_{\to m}\right) = \sum_{j=1}^{m} \gamma_j \mathbb{E}_{F_i|\mathbf{a}_{\to m},\mathbf{y}_{\to m}} \mathrm{L}_j\left(\mathbf{f}_j,\mathbf{a}_j\right) \tag{11}$$

Note that the linearity of the expectation allows it to be passed through the summation, and that, since the losses associated with each timestep are functions of the random variables associated with that timestep only, the expectation need only be taken over these variables. Thus, while the number of expectations is increased, the dimensionality of each expectation is decreased significantly.

The general value function is also expressed differently. It is now interpreted as a 'cost-to-go' function, i.e. the total expected cost to manage the system from a given timestep forward, with the slight notational change from $\mathbb{E}\mathrm{L}^*_j$ to $\mathbb{E}\mathrm{L}^*_{j\to}$ to reflect that it is now the 'cost-to-go' rather than the total expected cost. The recursive definition from Eq. (4) now becomes:

$$\mathbb{E}\mathrm{L}^*_{j\to}\left(\mathbf{a}_{\to(j-1)},\mathbf{y}_{\to(j-1)}\right) = \mathbb{E}_{Y_j|\mathbf{a}_{\to(j-1)},\mathbf{y}_{\to(j-1)}} \min_{\mathbf{a}_j\in\mathcal{A}_j} \mathbb{E}\mathrm{L}_{j\to}\left(\mathbf{a}_{\to j},\mathbf{y}_{\to j}\right) \tag{12}$$

where:

$$\mathbb{E}\mathrm{L}_{j\to}\left(\mathbf{a}_{\to j},\mathbf{y}_{\to j}\right) = \mathbb{E}_{F_j|\mathbf{a}_{\to(j-1)},\mathbf{y}_{\to j}} \mathrm{L}_j\left(\mathbf{f}_j,\mathbf{a}_j\right) + \frac{\gamma_{j+1}}{\gamma_j}\mathbb{E}\mathrm{L}^*_{(j+1)\to}\left(\mathbf{a}_{\to j},\mathbf{y}_{\to j}\right) \tag{13}$$

That is, the cost to optimally manage the system from timestep j onward is the expectation (over the new measures which first become available at timestep j) of the minimum (over the actions selected at timestep j) of the expected loss at timestep j plus the discounted loss of managing the system from timestep $j+1$ onward. Note that $\frac{\gamma_{j+1}}{\gamma_j}$ represents the factor for discounting the loss at timestep $j+1$ back to timestep j.

The posterior expected loss, similar to Eq. (6), is:

$$\mathbb{E}\mathrm{L}(\mathrm{Y}) = \gamma_1\mathbb{E}\mathrm{L}^*_{1\to}\left(\mathbf{a}_{\to 0},\mathbf{y}_{\to 0}\right) = \mathbb{E}_{\mathbf{Y}_1} \min_{\mathbf{a}_1\in\mathcal{A}_1} \left[\mathbb{E}_{\mathbf{F}_1|\mathbf{y}_1}\gamma_1\mathrm{L}_1\left(\mathbf{f}_1,\mathbf{a}_1\right) + \gamma_2\mathbb{E}\mathrm{L}^*_{2\to}\left(\mathbf{a}_{\to 1},\mathbf{y}_{\to 1}\right)\right] \tag{14}$$

The prior expected loss is evaluated similarly by dropping any conditioning on observations from Eqs. (12–14), and VoI is again evaluated as in Eq. (9).

Evaluation of VoI in systems with decomposable loss functions involves a lower problem dimensionality with respect to the evaluation of VoI under general loss functions. Under the assumption of decomposability, the bottleneck is the need to evaluate $\mathbb{E}\mathrm{L}^*_{m\to}(\mathbf{a}_{\to(m-1)},\mathbf{y}_{\to(m-1)})$ for each possible sequence of past actions and measurements. Returning to the reference case of discrete variables and actions, this dimensionality is quantified as $\mathcal{O}(m\theta_y^{|Y|}\theta_a^{|\mathcal{A}|}\theta_f^n)$, i.e., for each of the $\theta_y^{|Y|}\theta_a^{|\mathcal{A}|}$ possible sequences of measurements and actions, an expectation must be taken over the θ_f^n possible states of random variables in the final timestep. Linear dependence on m results from the need to repeat a similar evaluation for each timestep, working backwards from j=m to j=1. Assuming that variables are binary and that $|Y|$ and $|\mathcal{A}|$ grow linearly with time, the overall dimensionality for the computation can be expressed as $\mathcal{O}(m2^{m(n_y+n_a)+n})$. Note that this dimensionality, while still growing exponentially in the management time duration, will do so at a slower rate than in the case of a general loss function, especially considering that typically the number of action choices and measurements are smaller than the number of random variables governing system performance in each timestep, i.e. $n_y \ll n$ and $n_a \ll n$.

Finally, it should be noted that POMDPs represent a special case of systems with decomposable loss functions discussed here. First, in a POMDP, it is assumed that the loss function associated with each timestep is of the same form, and therefore that the same sets of possible actions and random variables are present for each timestep. Second, the key assumption of a Markov process is that the current joint state of the random variables affecting the system is independent of all past states given the joint state of the previous timestep and any actions taken and observations made between these steps. For this reason, POMDP solution methods typically make use of a belief state which is the posterior probability distribution over the state variables at the current timestep, $\mathrm{p}_{F_j|\mathbf{a}_{\to j},\mathbf{y}_{\to j}}$. This belief state represents a sufficient statistic describing the current and future state of the system, since the belief state at timestep $j+1$ can be obtained from that at timestep j given actions $\mathbf{a}_j$ and observations $\mathbf{y}_j$ made between these timesteps. This contrasts with the general case, where $\mathrm{p}_{F_{j+1}|\mathbf{a}_{\to j+1},\mathbf{y}_{\to j+1}}$ must be obtained from prior distribution $\mathrm{p}_{F_{j+1}}$ together with all actions $\mathbf{a}_{\to j+1}$ and all observations $\mathbf{y}_{\to j+1}$. In POMDPs, the value function is defined as a function of the belief state rather than of all past observations and actions to take advantage of the Markovian property. For a POMDP, if $\mathscr{b}_j$ denotes the belief state at timestep $\mathscr{b}_j$, then Eqs. (12) and (13) are expressed together as:

$$\mathbb{E}\mathrm{L}^*_{j\to}(\mathscr{b}_j) = \mathbb{E}_{Y_j|\mathscr{b}_j} \min_{\mathbf{a}_j\in\mathcal{A}_j} \left[\mathbb{E}_{F_j|\mathscr{b}_j,\mathbf{y}_j}\mathrm{L}_j(\mathbf{f}_j,\mathbf{a}_j) + \frac{\gamma_{j+1}}{\gamma_j}\mathbb{E}\mathrm{L}^*_{(j+1)\to}(\mathscr{b}_{(j+1)})\right] \tag{15}$$

where $\mathscr{b}_{(j+1)}$ is a function of $\mathscr{b}_j$, $\mathbf{a}_j$, and $\mathbf{y}_j$. This is an example of the classical Bellman Equation used for the solution of POMDP problems [22]. In the case of discrete variables, $\mathscr{b}_j$ would be a vector encoding the probabilities of all joint states of the variables at timestep j. Also note that in some Bellman Equation formulations, the expectation over observations at timestep j, which is the outermost operation performed here, can be moved within, forming the innermost expression, i.e., as an expectation over observations at timestep $j+1$ of $\mathbb{E}\mathrm{L}^*_{(j+1)\to}(\mathscr{b}_{(j+1)})$, where under this formulation $\mathscr{b}_{(j+1)}$ would be a function of $\mathbf{y}_{(j+1)}$ rather than $\mathbf{y}_j$.

2.4. Uncontrolled system evolution

We now consider that, in addition to the system having a decomposable loss as in Eq. (10), that the actions taken to manage the system have no effect on the evolution of the random variables which affect the system. That is, $\mathrm{p}_{F|\mathbf{a}} = \mathrm{p}_F \forall\ \mathbf{a} \in \mathcal{A}$. Thus, the influence of actions on the system is felt only through the decomposable loss function. In other words, while the agent cannot control the evolution of the system directly, he or she can control the impact which the system state will have on the system management cost for each timestep by choosing an appropriate response action. Recall that the example problem of Section 1 had such a form, where decisions made in response to the random field state reduced the penalties incurred without changing the field itself. In general, any system where the effects of management actions are confined to a single timestep only, or are relatively limited in scope compared to the management lifetime, might be effectively modeled as an uncontrolled system. Preparation for extreme events is a natural application, since emergency precautions taken to protect the system will not alter the underlying mechanisms by which the extreme event process occurs. Although this assumption is quite restrictive in terms of the types of system management activities which can be modeled, it is also, as demonstrated below, quite powerful in terms of reducing the overall computational complexity of VoI evaluation. Therefore, in situations

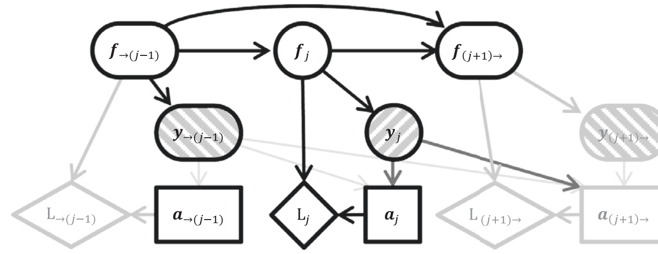

Fig. 5. Probabilistic graphical model for an uncontrolled system with a decomposable loss function. Potential interdependencies between measurements have been removed to improve clarity.

where the system of interest can be appropriately modeled as an uncontrolled system, it is computationally of great advantage to make use this assumption.

Fig. 5 depicts a PGM for such an uncontrolled system. Note that actions now only influence the loss associated with each timestep, rather than the evolution of the underlying random variables as in previous cases.

In an uncontrolled system, losses can be again evaluated via the dynamic programming approach discussed in Section 2.3. However, conditioning of current and future measurements and random field variable states on past actions can be ignored following the assumption of uncontrollability. This allows Eqs. (12) and (13) to be expressed together as:

$$\mathbb{E}\mathrm{L}^*_{j\to}(\mathbf{y}_{\to(j-1)}) = \mathbb{E}_{Y_j|\mathbf{y}_{\to(j-1)}} \min_{\mathbf{a}_j\in\mathcal{A}_j}\left[\mathbb{E}_{F_j|\mathbf{y}_{\to j}}\mathrm{L}_j(\mathbf{f}_j,\mathbf{a}_j)\right] + \frac{\gamma_{j+1}}{\gamma_j}\mathbb{E}_{Y_j|\mathbf{y}_{\to(j-1)}}\mathbb{E}\mathrm{L}^*_{(j+1)\to}(\mathbf{y}_{\to j}) \tag{16}$$

This provides a recursive definition for $\mathbb{E}\mathrm{L}^*_{j\to}(\mathbf{y}_{\to(j-1)})$, the value function in the case of an uncontrolled system. Note that this is only a function of past measurements; within Eq. (16), the choice of actions for timestep j is optimized, and since expectations over current measurements and random field variable states are not a function of past actions, dependence of the function on these actions is removed.

The posterior expected loss under measurement scheme Y is evaluated as in Eq. (14):

$$\mathbb{E}\mathrm{L}(\boldsymbol{Y}) = \gamma_1\mathbb{E}\mathrm{L}^*_{1\to}(\mathbf{y}_{\to 0}) \tag{17}$$

Note that if we substitute in for the value functions using Eq. (16), distributing the discounting factors and measurement expectations across the summands and using the chain rule to collect these expectations, we obtain the following closed-form expression for the posterior expected loss for an uncontrolled system with a decomposable loss function:

$$\mathbb{E}\mathrm{L}(\boldsymbol{Y}) = \sum_{j=1}^{m}\gamma_j\mathbb{E}_{Y_{\to j}}\min_{\mathbf{a}_j\in\mathcal{A}_j}\left[\mathbb{E}_{F_j|\mathbf{y}_{\to j}}\mathrm{L}_j(\mathbf{f}_j,\mathbf{a}_j)\right] \tag{18}$$

Similarly, the prior expected loss can be evaluated in closed form as:

$$\mathbb{E}\mathrm{L}(\emptyset) = \sum_{j=1}^{m}\gamma_j\min_{\mathbf{a}_j\in\mathcal{A}_j}\left[\mathbb{E}_{F_j}\mathrm{L}_j(\mathbf{f}_j,\mathbf{a}_j)\right] \tag{19}$$

VoI is again evaluated as in Eq. (9).

VoI computation in uncontrolled systems involves evaluating each term of Eqs. (18) and (19) separately and summing the results. For the entire evaluation, the problem dimensionality referring to the discrete reference problem is $\mathcal{O}(m\theta_y^{|Y|}\theta_a^{|\mathcal{A}_m|}\theta_f^m)$. Assuming a linear growth of binary measurement and action set sizes with the time duration, this dimensionality is expressed as $\mathcal{O}(m2^{mn_y+n_a+n})$. Again, while the growth in dimensionality remains exponential in the management time duration, the rate of growth is significantly reduced compared to the previous cases, especially considering that the number of observations taken in each timestep is typically small compared to the number of random variables or management actions, i.e. $n_y \ll n$ and $n_y \ll n_a$. Furthermore, if the number of measurements is independent of the management time duration, e.g., if problem constraints dictate a fixed upper limit to the number of measures obtained, any exponential dependence between problem dimensionality and the management time duration is removed. Thus, under the uncontrolled system assumption, holding all other parameters (i.e. the number of possible states and actions per timestep and the total number of observations of the system) fixed, the growth of computational complexity for evaluation of the VoI is linear in the time duration length. Also note that if the loss function is decomposable in space as well as in time, as in Eq. (32), then the complexity will be linear in the number of system components [10,14].

Finally, mention should be made of the special case of loss quantified via the L-2 norm of the prediction error, where the actions are interpreted as "guesses" on the states of the random variables and the loss is proportional to the sum of the squared error between these guesses and the true field values. Note that this is a case of an uncontrolled system, as the guesses have no impact on the behavior of the random field. As discussed by Malings and Pozzi [10] for spatial systems, the VoI under this loss function is a function of the prior and posterior covariance of the random field. In Gaussian random fields, this covariance is not a function of the measurement values, but only of their locations. Similarly, for spatio-temporal systems described by Gaussian random fields, the VoI under this loss function is only a function of the sensor placement and scheduling scheme, and not of the specific outcomes of measurements. Therefore, sensor placement and scheduling can be performed using the VoI metric without the need to take the expectation over potential measurement outcomes.

3. Sensor placement and scheduling

The VoI metric can be used to identify an optimal sensing scheme Y^*, i.e. a set of measurement locations and times for the system, from a set $\mathcal{Y}$ of potential measurement locations and times as follows:

$$Y^* = \text{argmax}_{Y \subseteq \mathcal{Y}} \ \text{VoI}(Y) - \text{C}(Y) \text{ subject to } \text{C}(Y) \le b \tag{20}$$

where $\text{C}(Y)$ represents the cost of implementing measurement scheme Y and b represents a fixed measurement budget. Appropriate definitions of the cost function can be used to impose constraints relevant to sensor placement and/or scheduling. For example, in a sensor placement problem, subsequent measures taken at a location which has been measured previously may not incur any additional cost, or may only incur a negligible cost. In this way, the low additional costs of interrogating a sensor which has already been installed at a certain location can be captured. Furthermore, in a problem where inspections of the system are carried out via an onsite agent such as a human or robotic inspector, the cost of additional measurements collected at the same time may be rather small, while the cost of additional measurements at a time when no other measures are scheduled is large. This captures the high cost of deploying the inspector to the site, but the low subsequent costs of collecting additional information once an inspector is already present.

Eq. (20) represents a problem in combinatorial optimization. In general, the only approach to combinatorial optimization which guarantees an optimal solution is exhaustive enumeration, i.e. computing the VoI of each possible subset of $\mathcal{Y}$. This is generally not a feasible solution approach, since the number of potential subsets grows exponentially with the size of $\mathcal{Y}$. To avoid this computational difficulty in this paper, we make use of an approximate greedy optimization approach, as described in Algorithm 1. This approach, based on previous work in greedy optimization e.g. [32], iteratively builds the optimal sensing set by adding single measurements to the set which most improve the objective function. Unfortunately, this algorithm does not guarantee optimal solutions to the objective in all cases. However, even when complete enumeration is infeasible, greedy optimization often leads to near-optimal results in many practical applications e.g. [14,10]. Further details on the greedy algorithm, its performance, and alternative approaches are discussed in the companion paper [20].

Algorithm 1 : Pseudo-code for the forward greedy algorithm.

Input candidate set $\mathcal{Y}$, objective function $\text{VoI}(\cdot)$, cost function $\text{C}(\cdot)$, budget b
$k = 0$, $Y_0 = \emptyset$
while $|\mathcal{Y}| > 0$
  $k = k + 1$
  select $y_k^* = \text{argmax}_{y \in \mathcal{Y}} \text{VoI}(Y_{k-1} \cup \{y\}) - \text{C}(Y_{k-1} \cup \{y\})$
  if $\text{C}(Y_{k-1} \cup \{y_k^*\}) \le b$
    $Y_k = Y_{k-1} \cup \{y_k^*\}$
  end
  $\mathcal{Y} = \mathcal{Y} \setminus \{y_k^*\}$
end
$k_{\text{end}} = k$
Output $Y^* = \text{argmax}_{Y \in \{Y_0, Y_1, \ldots, Y_{k_{\text{end}}}\}} \text{VoI}(Y) - \text{C}(Y)$

Algorithm 1 represents an approach to solving the offline optimal sensor placement and scheduling problem. In this offline problem, the optimal sensing scheme is decided upon before it is implemented, and is implemented according to the pre-selected scheme. However, as measurements are collected under a given offline scheme, additional information about the system may be used to revise this scheme for future timesteps. This is referred to as the online optimal sensing problem, in which the sensing scheme is re-evaluated and re-optimized at each timestep to reflect the latest knowledge of the system. Online sensing allows for a greater flexibility, as the information collection plan is updated to conform to the current state of knowledge of the system. For this reason, in an expected sense, the online sensing approach will outperform (i.e., perform at least as well as) the offline approach in terms of the losses incurred for system management. Again, while guarantees on greedy optimization performance are available for sensing metrics, these do not hold in general for the VoI [33,34].

To perform online sensor placement optimization based on the VoI metric, at each timestep, a revised optimal sensing plan should be determined which selects the best plan for future observations based on data collected by past measurements as well as past actions implemented to manage the system based on these observations. Solutions to these sensing problems may be obtained using suitable approaches to combinatorial optimization, e.g. the forward greedy optimization approach of Algorithm 1. In the online case, the optimal set of future observations is based on maximization of the net marginal VoI. The marginal VoI, i.e. the additional benefit of obtaining future measurements $Y_{j\to}$ given that past measures $\mathbf{y}_{\to(j-1)}$ have already been collected and that actions $\mathbf{a}^*_{\to(j-1)}$ have already been implemented, is evaluated as:

$$\begin{aligned}&\text{VoI}\left(Y_{j\to}|\mathbf{a}^*_{\to(j-1)}, \mathbf{y}_{\to(j-1)}\right)\\&= \mathbb{E}\text{L}\left(\emptyset|\mathbf{a}^*_{\to(j-1)}, \mathbf{y}_{\to(j-1)}\right) - \mathbb{E}\text{L}\left(Y_{j\to}|\mathbf{a}^*_{\to(j-1)}, \mathbf{y}_{\to(j-1)}\right)\end{aligned} \tag{21}$$

where the marginal loss is evaluated in general, following Eq. (6), as:

$$\mathbb{E}\text{L}\left(Y|\mathbf{a}^*_{\to(j-1)}, \mathbf{y}_{\to(j-1)}\right) = \mathbb{E}\text{L}^*_j\left(\mathbf{a}^*_{\to(j-1)}, \mathbf{y}_{\to(j-1)}\right) \tag{22}$$

Note that in systems with decomposable loss functions, the efficient computational approaches of Sections 2.3 (for controlled systems) and 2.4 (for uncontrolled systems) can also be applied to evaluate marginal VoI. Also note that online optimization is computationally more challenging than offline optimization due to the need to re-optimize the selection of the sensing scheme for future measurements at each timestep. Thus, the computational complexity of online optimization is greater than that of offline optimization by a factor of m. Furthermore, online optimization will be impossible when the computational time required for the optimization is longer than the timestep duration. Also note that the marginal VoI should be traded off against the total cost of all measurements $\text{C}(Y^*_{\to(j-1)} \cup Y_{j\to})$. However, the budget constraint of Eq. (20) might either be applied to the total cost of all measurements or to the cost for measurements in the next timestep only, i.e., $\text{C}(Y_j) \le b_j$.

In the latter case, the constraint in Algorithm 1 should be appropriately modified. Finally, note that for the special case of the L-2 norm of the prediction error in Gaussian random fields, there is no benefit to online sensor placement, as the VoI based on this loss is merely a function of the prior and posterior covariance of the random field, which are not affected by the specific outcomes of measurements. For general loss functions, however, different measurement outcomes will lead different future measurement schemes to be optimal, e.g. as symptoms of impending future events are detected (as in Fig. 1e–f).

4. Gaussian random field models

In the examples which follow, Gaussian random field models are used to define the distributions of the random variables affecting systems. It should be noted that the approaches for VoI evaluation presented in the preceding sections do not require the use of such a model; any appropriately structured PGM, such as a dynamic Bayesian network, might be used. However, Gaussian random field models are one of the few model types which allow for efficient closed-form Bayesian updating, as in Eq. (26) below. They are defined completey by their mean and covariance structure. They also provide further opportunities for efficiencies in the computation of VoI, as discussed in previous work [14]. For alternative forms of PGM, approximate inference techniques including Markov Chain Monte Carlo methods or particle filters may be used, incurring a higher computational cost and/or risking lack of convergence of the numerical procedure [29,35].

Gaussian random field models represent a generalization of the multivariate Gaussian distribution to a continuous domain, and have been used to represent a wide variety of spatio-temporal phenomena [36]. A similar model is used by Malings and Pozzi [10] to describe purely spatial systems, and is extended to spatio-temporal systems here as:

$$\mathrm{f}(\mathbf{x},t) \sim \mathcal{GP}\left[\mu(\mathbf{x},t), \mathrm{k}\left(\mathbf{x},t,\mathbf{x}',t'\right)\right] \tag{23}$$

where mean function $\mu(\mathbf{x}, t)$ describes the mean of the random field at spatial location $\mathbf{x}$ and time t, and covariance function $\mathrm{k}(\mathbf{x}, t, \mathbf{x}', t')$ defines the covariance between random variables at $\mathbf{x}$; t and $\mathbf{x}'$; t'. Over any finite discretized spatio-temporal domain, this model defines a multivariate Gaussian distribution for the random variables affecting the system over this domain:

$$\mathbf{f} \sim \mathcal{N}\left(\boldsymbol{\mu}_F, \boldsymbol{\Sigma}_F\right) \tag{24}$$

where mean vector $\boldsymbol{\mu}_F$ and covariance matrix $\boldsymbol{\Sigma}_F$ are derived by evaluating the mean and covariance functions at all spatio-temporal coordinates and combinations of coordinates in the domain.

By observing these random variables, or linear combinations of these variables, a Gaussian observation vector is defined as follows:

$$\mathbf{y} = \mathbf{R}_Y\mathbf{f} + \boldsymbol{\epsilon} \qquad \boldsymbol{\epsilon} \sim \mathcal{N}\left(\boldsymbol{\mu}_\epsilon, \boldsymbol{\Sigma}_\epsilon\right) \tag{25}$$

where observation matrix $\mathbf{R}_Y$ encodes relationships between measurements within scheme Y and the random variables (or linear combinations of variables) which are observed. For example, if this matrix is used to encode two measurements, the first of which is a measurement of the third element of $\mathbf{f}$ and the second of which is a measurement of the average of the first three elements of $\mathbf{f}$, the observation matrix would have two rows of the same length as $\mathbf{f}$, with the first row having an entry of 1 in the third position and the second row having entries of 1/3 in the first, second, and third positions. In general, the observation matrix should be defined appropriately to encode the relationships between various physical sensors deployed in the system and the quantities which they are measuring in any given application, where each row of $\mathbf{R}_Y$ corresponds to a measurement in scheme Y and each column corresponds to a variable in random field vector $\mathbf{f}$. Furthermore, temporal constraints must be obeyed, i.e. measurements associated with a specific time should not be dependent on random variables associated with future times. Measurement noise is encoded in the noise vector ϵ, which is assumed to have a multivariate Gaussian distribution with mean $\boldsymbol{\mu}_\epsilon$ and covariance $\boldsymbol{\Sigma}_\epsilon$. This noise may in general have a bias via a non-zero mean vector or may be correlated between measurements at different locations or times via appropriate definition of the covariance. Based on the definition of Eq. (25), the vector $\mathbf{y}$ of measurements itself has a multivariate Gaussian distribution, with mean vector $\boldsymbol{\mu}_Y = \mathbf{R}_Y\boldsymbol{\mu}_F + \boldsymbol{\mu}_\epsilon$ and covariance matrix $\boldsymbol{\Sigma}_Y = \mathbf{R}_Y\boldsymbol{\Sigma}_F\mathbf{R}_Y^{\mathrm{T}} + \boldsymbol{\Sigma}_\epsilon$. Given a vector of measures $\mathbf{y}$, the prior Gaussian model of the random variables affecting the system can be updated to a posterior model:

$$\mathbf{f}|\mathbf{y} \sim \mathcal{N}\left(\boldsymbol{\mu}_{F|\mathbf{y}}, \boldsymbol{\Sigma}_{F|Y}\right) \tag{26}$$

where $\boldsymbol{\mu}_{F|\mathbf{y}} = \boldsymbol{\mu}_F + \mathbf{R}_Y\boldsymbol{\Sigma}_F\boldsymbol{\Sigma}_Y^{-1}(\mathbf{y} - \boldsymbol{\mu}_Y)$ and $\boldsymbol{\Sigma}_{F|Y} = \boldsymbol{\Sigma}_F - \mathbf{R}_Y\boldsymbol{\Sigma}_F\boldsymbol{\Sigma}_Y^{-1}\boldsymbol{\Sigma}_F\mathbf{R}_Y^{\mathrm{T}}$.

5. Case study applications

This section presents two case studies illustrating how the VoI metric can be applied to determine optimal sensor placement and scheduling schemes for infrastructure systems. In Section 5.1, a simulated system of columns subjected to differential settlement is investigated, and offline and online optimal sensing schemes for the system are evaluated and compared. In Section 5.2, optimal sensor placement and scheduling schemes are determined making use of data collected for the Sherman and Joyce Bowie Scott Hall building, a recently constructed building at Carnegie Mellon University whose main structural elements have been instrumented with strain sensors.

5.1. Application to differential settlement

An example problem is presented here to illustrate the application of the VoI evaluation and optimal sensor placement and scheduling methods outlined above. This example is motivated by the monitoring of and response to settlement under columns of a structure over time [37]. The settlements under $n = 9$ columns of a structure over an $m = 10$ year period are modelled by a random field using a Gaussian random field. The physical arrangement of the columns is identified in Fig. 7a, while the time duration considered is discretized as $T = \{1, 2, \dots, 10\}$ years. The mean function of the Gaussian random field is:

$$\mu(\mathbf{x},t) = \mu_0\left[1 - \exp\left(-\frac{t-t_0}{\alpha_t}\right)\right] \tag{27}$$

where, $\mu_0 = 0.5$ m, $t_0 = 1$ year, and $\alpha_t = 5$ years. This mean function models the settlement of the columns over time, with the average amount of settlement of the columns increasing to a long-term average of μ_0. The covariance function is defined to be decomposable between space and time, as follows:

$$\mathrm{k}\left(\mathbf{x},t,\mathbf{x}',t'\right) = \sigma(t)\sigma\left(t'\right)\rho_X\left(\mathbf{x},\mathbf{x}'\right)\rho_T\left(t,t'\right) \tag{28}$$

The spatial component is:

$$\rho_X\left(\mathbf{x},\mathbf{x}'\right) = \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}'\|^2}{2\lambda_X^2}\right) \tag{29}$$

This model represents a square exponential correlation function, where the settlements of nearby columns are more heavily correlated. The range of the correlation is parameterized by the correlation length $\lambda_X = 20$ m. This spatial correlation structure imposes certain relationships on the relative settlements of the columns in the structure at a specific time. The temporal component is:

$$\rho_T\left(t,t'\right) = \exp\left(-\frac{\left|t-t'\right|^2}{2{\lambda_T}^2}\right) \tag{30}$$

This temporal component again models a square exponential correlation between settlements at different times, with a correlation timescale of $\lambda_T = 5$ years. This governs the differentiability or smoothness of the

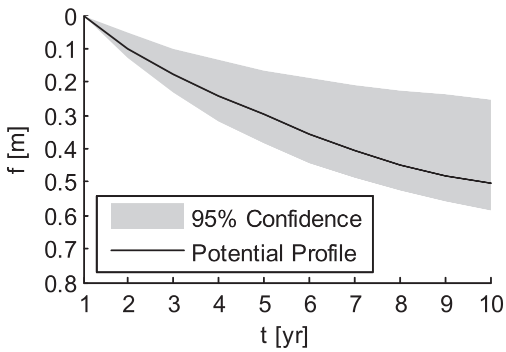

Fig. 6. Prior 95% confidence region for settlement of a column over time (grey area) with a potential settlement profile indicated (black line).

column settlement profiles in time, ensuring infinite differentiability. Finally, the standard deviation function for the settlement is:

$$\sigma(t) = \sigma_0\left[1 - \exp\left(-\frac{t - t_0}{\alpha_t}\right)\right] \tag{31}$$

where $\sigma_0 = 0.1$ m. This models an increase in the variance of the settlement over time where, comparing Eqs. (27) and (31), the coefficient of variation of the settlement at any time is a constant, i.e., $\mu_0/\sigma_0 = 5$.

It is important to note that the forms and relevant parameters of Eqs. (27–31) are intended to represent a possible reasonable probabilistic model for structural settlement. The selection of this model, including the assumption of independence between spatial and temporal covariance of Eq. (28), choice of correlation function forms, and the setting of parameters, should all be performed using available prior knowledge. The reader is referred elsewhere for discussions of methods and examples of the creation and calibration of probabilistic spatio-temporal models [6,8,36].

Based on this prior mean and variance model, the prior 95% confidence region for settlement of a column over time, as well as a potential realization of a settlement profile, are depicted in Fig. 6. Note the decrease in the average settlement and increase in uncertainty over time.

Within this system, observations of the column settlements are possible for each column annually. This would define the observation matrix $\mathbf{R}_{\mathcal{Y}}$ for the set of all possible measurements to be an identity matrix of size 90 (corresponding to measures of each of 9 columns in each of 10 years). Observation matrices $\mathbf{R}_Y$ for particular measurement schemes Y can be obtained from $\mathbf{R}_{\mathcal{Y}}$ by eliminating rows of the matrix corresponding to potential observations in $\mathcal{Y}$ which are not included in scheme Y. Errors in the settlement measurements are considered for this problem. Measurement errors are modeled as independent Gaussian random variables with a mean of zero and a standard deviation of $\sigma_\epsilon = 0.01$ m. Therefore, $\boldsymbol{\mu}_\epsilon = \mathbf{0}$ and $\boldsymbol{\Sigma}_\epsilon = \sigma_\epsilon^2\mathbf{I}$, where $\mathbf{0}$ represents a zero vector and $\mathbf{I}$ represents an identity matrix of an appropriate size. The impacts of correlated errors in settlement measurements are discussed in related work [38]. The outcomes of these measures are assumed to not be available to the managing agent until the following timestep, i.e. decisions made in year j can be based on gathered information about the settlement of the columns in years 1 through $j-1$.

To manage this system, in year j, the managing agent has an option to intervene by selecting $\mathbf{a}_{j,i} = 1$ at cost C_r=\$10k, to prevent damage due to excessive settlement of column i compared to the average settlement of all columns, which would incur a cost C_f=\$100k. Otherwise, the agent would choose to do nothing, i.e. select $\mathbf{a}_{j,i} = 0$. To allow for the assumption of an uncontrolled system, it is important that the modeled intervention activity not affect the settlement of the columns themselves. Thus, an intervention to search for and repair damage due to settlement, e.g. patching façade cracks and re-squaring door frames, which mitigates damages due to differential settlement without affecting the settlement itself, is considered here. This problem is encoded by

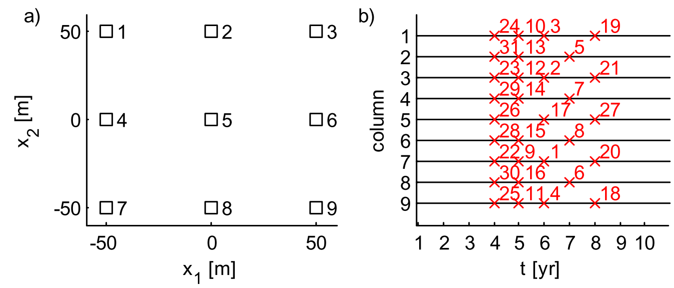

Fig. 7. (a) Spatial arrangement of the structural columns; (b) Measurement times and locations for the scheme optimized offline based on VoI (x). Numbers indicate the order in which measures are selected by the greedy algorithm.

the following annual loss function, which is decomposable in space as well as time:

$$\mathrm{L}_j(\mathbf{f}_j, \mathbf{a}_j) = \sum_{i=1}^{n} \mathrm{L}_{j,i}(\mathbf{f}_{j,i}, \mathbf{a}_{j,i}) \tag{32}$$

where:

$$\mathrm{L}_{j,i}(\mathbf{f}_{j,i}, \mathbf{a}_{j,i}) = \begin{cases} C_f \mathbb{I}\left(\left|\mathbf{f}_{j,i} - \frac{1}{n}\sum_{i=1}^{n}\mathbf{f}_{j,i}\right| > \delta_{\max}\right) & \text{if } \mathbf{a}_{j,i} = 0 \\ C_r & \text{if } \mathbf{a}_{j,i} = 1 \end{cases} \tag{33}$$

where $\mathbb{I}(\cdot)$ is an indicator function, taking on value 1 when its argument is true and 0 otherwise, and $\delta_{\max} = 0.1$ m. The total loss is the discounted cumulative loss over the time duration as in Eq. (10), with a discount factor of $\gamma_j = 0.9^{(j-1)}$. In this problem, sensing costs are assumed to be \$1k per measurement, discounted to the present value using the same discounting method.

Based on the probabilistic model and decision-making problem outlined above, an agent managing the system should plan on implementing intervention actions for columns later on during the management time duration. This is because, as time increases, the variances of the column settlements increase, as in Eq. (31), and therefore the likelihood that a column's settlement will differ from the average by more than the acceptable threshold $\delta_{\max}$ increases. Furthermore, the agent should plan on intervening earlier for the corner columns (1, 3, 7, and 9 in Fig. 7a) than for the side columns (2, 4, 6, and 8), and should intervene for the center column (5) last. This is because the spatial correlation structure of Eq. (29) means that the settlements of columns farthest from the spatial center of the domain are most likely to differ from the mean by more than the acceptable threshold. This intuition is reflected in the prior optimal management scheme depicted in Fig. 8a.

In terms of optimizing inspections based on the VoI, there are several factors which are considered and traded off. In time, earlier measurements of settlement will be available for use in supporting decision-making throughout the management time duration, but the absolute settlements of the columns are low early on, leading to low probabilities of settlements differing from the mean by more than the acceptable threshold. Later on, the absolute settlements of the columns are greater, leading to higher probabilities of failure and a higher chance that additional information will lead to better decision-making outcomes. However, while earlier measures can be used to support later decision-making, later measures cannot support earlier decisions, and so the value of these measures in supporting decisions at other timesteps is also diminished. In both space and time, the correlation structures mean that measurements which are close in space and/or time will tend to be redundant, and therefore appropriate spacing in space and time should be determined. Intuitively, an optimal sensing scheme would focus more measures toward the middle of the management time duration, trading off the factors discussed above. In space, this scheme would tend to stagger measurements between different columns, taking advantage of both spatial and temporal correlations to infer the behavior of unmeasured

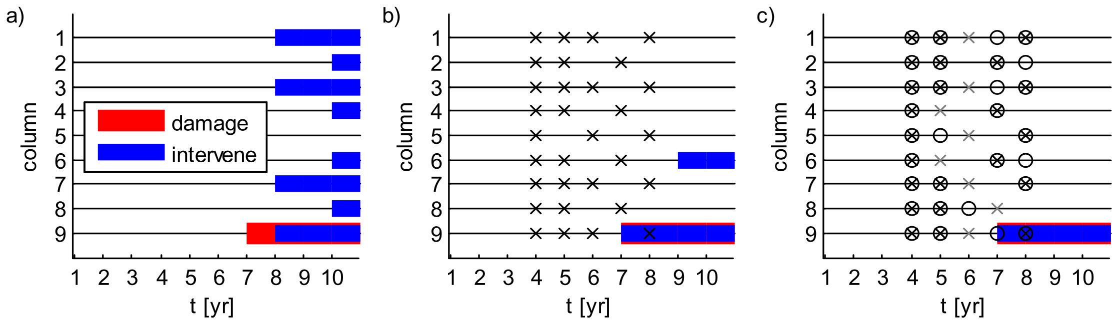

Fig. 8. (a) Prior optimal actions for managing the system and the state of the system for a specific realization of the system evolution; (b) Posterior optimal actions for system management based on the offline optimized sensing scheme (x) for the same realization; (c) Measurements selected via online optimization (o) and associated optimal actions.

Table 1
Comparison of the performance of various sensor placement and scheduling schemes.

	$\|Y\|$	$\mathbb{E}L(\emptyset)$	$\mathbb{E}L(Y)$	$VoI(Y)$	$C(Y)$	net VoI
No sensors	0	$170k	$170k	$0	$0	$0
Offline scheme	31	$170k	$32k	$138k	$19k	$119k
Intuitive scheme	45	$170k	$32k	$138k	$30k	$108k
All measures	90	$170k	$23k	$147k	$58k	$89k

Table 2
Comparison of the performance of online and offline sensor placement and scheduling schemes for a specific realization of the system's evolution, as depicted in Fig. 8.

	$\|Y\|$	$L(\mathbf{f}, \mathbf{a}^*)$	$L(\mathbf{f}, \mathbf{a}^*) + C(Y)$	net benefit
No sensors	0	$120k	$120k	$0k
Offline scheme	31	$27k	$46	$74k
Online scheme	30	18k	$37k	$83k

columns from earlier measurements on that column as well as from measures on nearby columns.

An offline optimal sensor placement and scheduling scheme to support system management using the VoI metric is identified in Fig. 7b following the greedy approach of Algorithm 1, and conforms well to the various intuitive factors discussed above. This scheme distributes measurements fairly evenly across all columns spatially, and concentrates measurements between years 4 and 8. Note that in years 4 and 5, measurements are prescribed for every column (except the center column in year 5). This provides two 'baseline' measurements of the settlement for most columns, allowing their true settlements to be accurately determined from this time. In years 6 through 8, columns are observed alternately, with the corner and center columns observed in years 6 and 8 and the side columns observed in year 7. Because of the smoothness imposed by the spatial and temporal covariance, these intermittent observations are enough to generate reasonably accurate posterior predictions for the column settlements in these years.

Without additional information, the prior expected loss for managing the system, $\mathbb{E}L(\emptyset)$, is $170k. With the offline optimal sensing scheme of Fig. 7b, the posterior expected loss, $\mathbb{E}L(Y)$, is $32k, and therefore $VoI(Y)$=$138k. Taking into account the cost of making these measurements, the net VoI, $VoI(Y) - C(Y)$, is $119k. Thus, by making these measurements to support decision-making, the overall management cost for the system is reduced by 70%, in an expected sense. Furthermore, the value of complete information, i.e. the VoI which would be obtained if all 90 potential measurements were implemented, is $147k. The optimal sensor set therefore achieves 94% of the value of complete information while including less than a third of the possible measurements. Finally, we can compare the VoI of the proposed scheme with that of an intuitive scheme, where measurements are taken for all columns every other year, starting in year 1. This is an example of the type of measurement scheme which might be prescribed for the system, with regularly scheduled inspections for all components on a fixed schedule. This scheme would provide roughly the same VoI as the optimized scheme, but does so at a higher sensing cost, leading to a lower net VoI for this scheme. The expected loss, cost, and VoI of these schemes are listed in Table 1 for easy comparison.

We also investigate the benefits of online optimal sensing for this system. Note that each online sensing scheme will be different based on the sequence of measurements collected during year 4; based on this information, measurement locations and schedules for subsequent years will be updated. To illustrate this, Fig. 8a depicts a specific realization of the system, where damage occurs on column 9 from year 7 onward. This figure also depicts the optimal prior management plan for the system, as discussed previously. The cost which is actually incurred by managing the system according to this prior plan for the given realization of the system's evolution is $120k.

If the offline optimized sensing scheme is used to support system management, more appropriate management actions can be taken, as depicted in Fig. 8b. Here, the prescribed interventions for column 9 correspond to the times where the column would have been damaged without these interventions, and therefore any failure costs are avoided. However, there is still an intervention prescribed for the neighboring column 6 in years 9 and 10 which does not correspond to an actual damage condition, but rather to a high risk of damage based on the collected information. The online sensor placement scheme, depicted in Fig. 8c, allows this unnecessary intervention to be avoided. This is done through more appropriately allocating measures to more accurately predict the state of column 6. As can be seen, many of the measures originally prescribed for year 6 are forgone, and instead more measurements are conducted during years 7 and 8, including for column 6. This allows for the probability of damage in this column to be reassessed, and as a result for the unnecessary intervention action to be avoided. Overall, adopting an online optimization approach reduced total cost (including system management and sensing costs) by 20% compared to the offline sensing scheme for this particular realization of the system. A summary of the costs for managing the system under various measurement schemes is provided in Table 2.

Comparing across different cases of online sensor placement, based on 100 random simulations of the settlement of columns and their measurement following different online schemes, in all cases the online optimal scheme consisted of fewer measures than the offline optimal scheme, and in 84 cases the losses incurred by the managing agent were reduced as well. The cost reduction of online sensing compared to offline sensing ranged from 50% to −49% in these cases, with an average of 15.4%. The empirical cumulative distribution for the cost reduction, $\hat{F}(\Delta C)$, is shown in Fig. 9. Note that the cases of negative cost reductions (i.e. cost increases) result from the fact that, occasionally, the new online scheme fails to capture certain relevant information which would have been captured by the offline scheme. Intuitively, while a measure

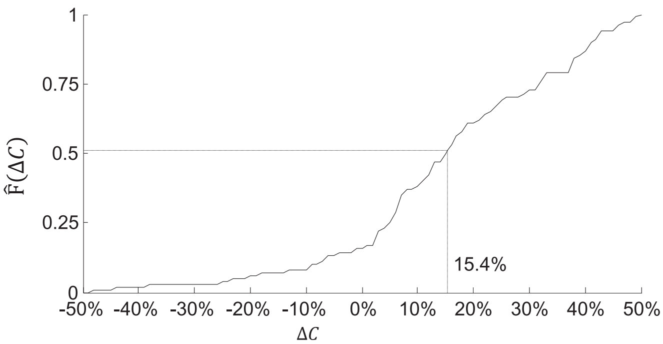

Fig. 9. Empirical cumulative distribution function ($\hat{\mathrm{F}}$) for the cost reduction (ΔC) from online sensing, with the average value indicated.

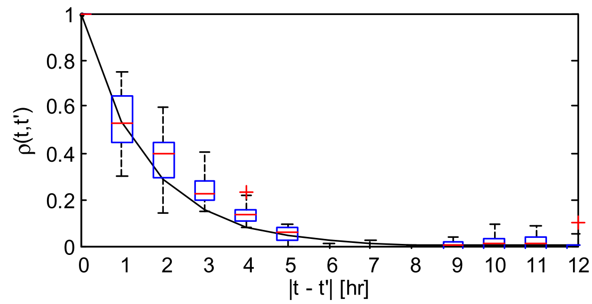

Fig. 10. Plot of empirical temporal correlations between column strain rates (boxplots represent ranges in empirical correlations across columns) versus the fitted temporal correlation model (black line).

may be neglected in an online scheme because previously collected information indicates that the column in question is likely 'safe', there is still a small but non-negligible chance that the column will become unsafe, and the denser set of offline measurements may have detected this. The relative benefit of online sensor placement compared to offline placement can thus vary greatly depending on the specific sequence of measurements obtained.

5.2. Example application for structural health monitoring

This section gives an example of the application of the above VoI evaluation framework to a problem of optimal sensor placement and scheduling for structural health monitoring using data obtained from an existing structure. The Sherman and Joyce Bowie Scott Hall is a newly constructed building on the campus of Carnegie Mellon University, housing a variety of academic and research spaces. A portion of the structure is suspended over a slope on eleven tubular steel columns. During construction, these columns were instrumented with 23 fiber optic strain gauges, as shown in Fig. 11a, with three sensors installed on column one and two sensors installed on all other columns. The physical arrangement of the columns is indicated in Fig. 11b. Columns are divided into three groups (indicated by colors in the diagram) which share common foundations.

In this problem, random field vector **f** represents the hourly strain rates in each of $n = 11$ columns over an $m = 24$ h period, representing a typical day during the construction process. These rates are described by a multivariate Gaussian distribution model. The mean hourly strain rate for each column was determined empirically from data collected during a seven-day model calibration period, which defined the function $\mu(\mathbf{x})$, the time-independent average column strain rate as a function of the position $\mathbf{x}$ occupied by the column within the structure. Similarly, the empirical spatial correlation between column strain rate data during this period was evaluated, and the spatial correlation function $\rho_X(\mathbf{x}, \mathbf{x}')$ was defined to be equal to this empirical correlation for any pair of columns at locations $\mathbf{x}$ and $\mathbf{x}'$. In the absence of empirical data, a finite element model of the structure might be used to generate simulated datasets for performing the calibration. Further details on the development and application of this spatial model can be found in previous work [39].

Additionally, the column strain rates were found to be temporally correlated. Analysis of data collected during the model calibration period in terms of the correlation of hourly column strain rates for individual columns as a function of the time difference of the data collection periods is presented in Fig. 10. To this empirical data, an exponential temporal correlation function of the following form was fitted:

$$\rho_T(t, t') = \exp\left(-\frac{|t - t'|}{\lambda_T}\right) \quad (34)$$

The correlation timescale parameter λ_T of this model was selected via maximization of the log-likelihood of the data set evaluated via a Gaussian random field model using this correlation function; a parameter value of $\lambda_T = 1.6$ hours was found to maximize the log-likelihood, and was therefore selected. A plot of the correlation function is provided as the black curve in Fig. 10.

The spatial and temporal correlation models were combined to define the full spatio-temporal model, following the assumption of decomposability as in Eq. (28). It should be noted that several assumptions, including that the mean column strain rate does not vary in time, the parametric form of the temporal correlation, that the spatial and temporal covariance of the column strain rates are decomposable, and that the column strain rates are well modeled as Gaussian random variables, have been incorporated into this model. This model for the behavior of Scott Hall appears to be reasonable based on the data collected, and is therefore used here for illustrative purposes. In the absence of this empirical data, data simulated from finite element models of the structure would have been used to develop this model.

Observations of the strain in the columns, as obtained via the strain sensors during a one-hour period, can be used to update this prior model of the strain rate, with an appropriate definition for $\mathbf{R}_Y$ used to define the relationship between the measured strain and the underlying random variables (the strain rate). For simplicity, a set of measurements obtained by a strain sensor associated with a particular column over an hour is pre-processed to determine the average measured strain rate in that column during that hour. This processed data is then considered to be a "measurement" of the column's average strain rate over that hour. Thus, rows in $\mathbf{R}_Y$ consist of zeros, except for a 1 in the position corresponding to the strain rate in vector **f** of the column in question at the hour when the sensor is active according to scheme Y. Observations are assumed to be independent with a nominal measurement error of 1% of the standard deviation of the measured quantity. This is based on the finding that the uncertainty in hourly strain rate conditional to strain gauge measurements was found to be several orders of magnitude smaller than the variability of the strain rates themselves.

We define a decision-making problem for system management based on the column strain rates. The loss function defining the problem decomposes temporally as in Eq. (10). For each timestep, the loss $\mathrm{L}_j(\mathbf{f}_j, \mathbf{a}_j)$ is determined as follows: if for any pair of column groups, the average strain rates of columns within these groups differ by more than a given threshold, then it its assumed that a harmful differential settlement between column foundations is occurring. This occurrence induces a loss of \$10k per hour while it persists, modeling the effort needed to correct such a settlement. However, the managing agent has the option to delay construction, at the cost of \$1k per hour, to avoid these consequences. No discounting is considered because of the short time period, i.e. $\gamma_j = 1\ \forall j \in 1, \ldots, m$. Decisions must be made an hour in advance, i.e. only the settlement measurements obtained via sensors up to and including the previous hour may be used to support decision-making. For this problem, the threshold of differential settlement is set to 4 mm per hour for illustrative purposes.

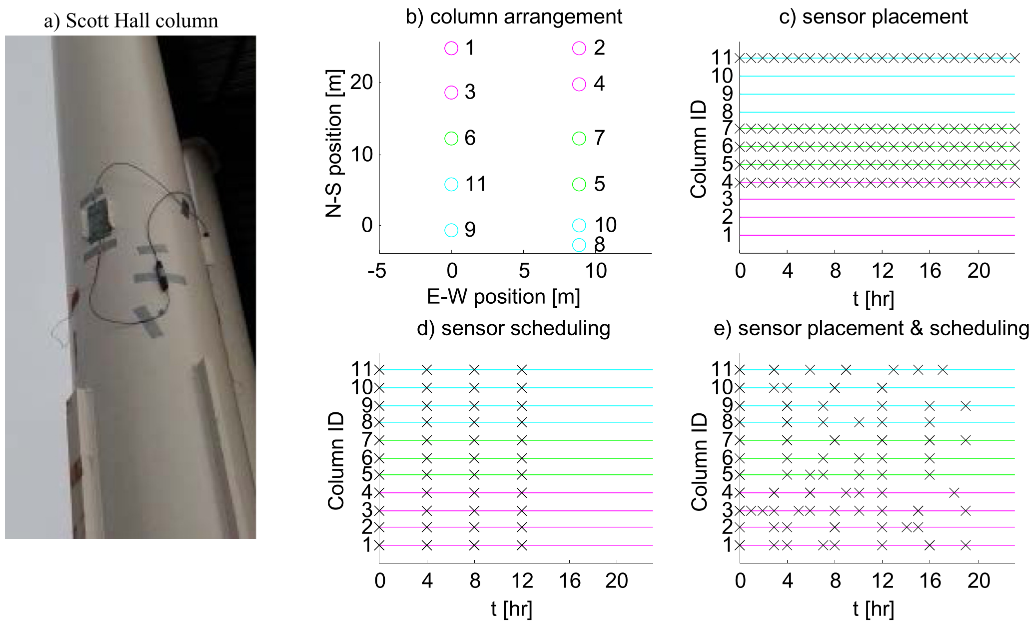

Fig. 11. (a) Image of a column of the Scott Hall building with attached strain sensor; (b) Spatial arrangement of the 11 instrumented columns of the structure, with colors indicating the column groups; (c) Sensor placement results; (d) Sensor scheduling results; (e) Simultaneous placement and scheduling results. The set of prescribed measurement times and locations in each case are indicated by x's in the spatio-temporal domain.

For the above decision-making problem, three alternative sensing schemes are optimized offline. First, sensor placement is considered, i.e. a subset of the full set of sensors actually installed on the structure is selected to gather data continuously to support decision-making for the above problem. Second, sensor scheduling is determined, i.e. certain times at which to measure strain rates for all available sensors are selected. Finally, simultaneous placement and scheduling are considered, where the times and locations at which to take measures are optimized.

Fig. 11c indicates results for optimal sensor placement. Results shown here correspond with the optimal spatial placement indicated in [39], consisting of five sensors spread across the three groupings of columns. Fig. 11d indicates results for optimal sensor scheduling. Here, some measurements are taken earlier in the cycle so that data will be available to guide later decision-making. Finally, Fig. 11e indicates results for optimal placement and scheduling. Here, prescribed measurements switch between columns as time passes, adjusting for changes in loading and strain patterns over time.

Overall, while the VoI of the sensor scheduling scheme is $ 1.2k, and that of the placement scheme is $ 2.8k, that of the optimal placement and scheduling is highest at $ 3.4k. While all schemes prescribe roughly the same number of sensor measurements, the optimal sensor placement and scheduling is the most flexible (as there are no constraints that all measurements on a particular column or at a particular time must be selected together), and therefore can provide greater benefits in terms of reduced management costs (as measured by VoI) while using a comparable number of sensors.

Note that, for this example, we investigate the performance of the method by selecting a threshold on strain rates which allows for the detection of periods of differential loading on the structure caused by the pouring of concrete during the construction process. These concrete pours are therefore used as proxies for failure events during actual building operations. Additionally, the costs associated with interventions and excessive strain rates are presented here for illustrative purposes.

6. Conclusions

The VoI metric is well suited to supporting optimal sensor placement and scheduling, as it directly identifies the benefits of information in reducing system management costs. However, this metric is difficult to evaluate in general due to the rapid growth in possible system states and management actions as the system management time duration increases. This paper identifies the special case of an uncontrolled system with a temporally decomposable loss function in which the VoI can be much more efficiently evaluated. The key assumptions for such a system are, first, that the loss function decomposes temporally, such that the total loss can be expressed as a discounted sum of losses across all timesteps, and second, that actions to manage the system do not influence the evolution of the random variables affecting the system for future timesteps. These assumptions can be relaxed somewhat through state augmentation, i.e., through extending the vector of actions and random variables associated with each timestep to include duplicates of variables associated with other timesteps. Such an approach has been applied previously for Markov Chain models, e.g. [40], and has allowed for the relaxation of certain model assumptions at the cost of increased computational complexity. In this way, the influences of short sequences of actions (e.g. a set of four coupled seasonal actions applied to a system in a year, which have no effect on the system in the next year) can be modeled, albeit at a higher computational cost. Additionally, previous work by Malings and Pozzi [10,41] has shown that more topologically complicated systems can, under certain circumstances, be evaluated under the cumulative system assumption (the spatial analog of the uncontrolled system assumption). It is possible that these insights can be extended to certain temporally evolving systems as well.

It should also be noted that, while the examples of Section 5 make use of binary action choices, any number of discrete action choices, or even continuous-valued actions, can also be handled. For discrete actions, the computational complexity will grow exponentially with the number of possible actions, unless the problem structure can be exploited to improve this (see [10] for examples). For continuous actions, problem structure might also be exploited in the case of convex problems, or surrogate models might be used to improve the efficiency of the optimization, e.g. [26]. Alternatively, the assumption of a discrete, finite time duration for the management problem is essential to the formulation of the efficient computational methods discussed here. In principle, depending on the discounting scheme, the influence of far-future actions and observations can be ignored, and an infinite-horizon problem can be approximated using a finite-horizon model. Furthermore, the effect of the time discretization has not been investigated; such a discretization might have a significant effect on the problem if the computation time is comparable to the length of the discretized timestep, in which case online optimization will not be possible following the approach presented here. New, more efficient, possibly heuristic techniques for online optimization would need to be implemented in that case.

In comparing sensor placement and scheduling in spatio-temporal systems based on VoI, as discussed in this paper, with problems of placement only, as discussed in previous work by Malings and Pozzi [10], several parallels can be found. Essentially, the temporal evolution of the system can be considered simply as another dimension, e.g., sensor placement and scheduling across a two-dimensional spatial domain can be seen as a special case of sensor placement in a three-dimensional domain, where two dimensions correspond to space and one to time. The key difference is in the evaluation of VoI in evolving systems, since information can propagate omnidirectionally in space, but it must propagate unidirectionally in time. In other words, while information collected at any spatial location in the system can support decision-making for any other spatial location, information collected at a specific time can only support decision-making for later times. However, evaluation of VoI in spatio-temporal systems via Eq. (18) can be seen as a special case of its evaluation in spatial systems, where only specific subsets of the full sensing scheme (corresponding to measurements associated with previous timesteps) are considered. Finally, online sensing in spatio-temporal systems can be compared to online sensing in purely spatial systems using batches of measurements (corresponding to measurements associated with a single timestep).

This paper also presents some examples to demonstrate the application of VoI evaluation in spatio-temporal systems and VoI-based sensor placement and scheduling. While schemes are optimized based on a forward greedy selection approach, both for online and offline placement and scheduling, as has been mentioned, there is no guarantee on the optimal outcome of this approach using the VoI metric. Further investigation of the optimization of measurement schemes based on the VoI metric, including shortcomings of the greedy approach illustrated here and potential approaches to overcoming these, are discussed in the companion paper [20].

Acknowledgments

The first author acknowledges the support of the Dowd Fellowship from the College of Engineering at Carnegie Mellon University. The authors would like to thank Philip and Marsha Dowd for their financial support and encouragement. The second author acknowledges the support of NSF project CMMI #1653716, titled “CAREER: Infrastructure Management under Model Uncertainty: Adaptive Sequential Learning and Decision Making”.

References

[1] Gomes WJS, Beck AT, Haukaas T. Optimal inspection planning for onshore pipelines subject to external corrosion. Reliab Eng Syst Saf Oct. 2013;118:18–27.
[2] Lam JYJ, Banjevic D. A myopic policy for optimal inspection scheduling for condition based maintenance. Reliab Eng Syst Saf Dec. 2015;144:1–11.
[3] Zhao X, Al-Khalifa KN, Nakagawa T. Approximate methods for optimal replacement, maintenance, and inspection policies. Reliab Eng Syst Saf Dec. 2015;144:68–73.
[4] Hajipour Y, Taghipour S. Non-periodic inspection optimization of multi-component and k-out-of-m systems. Reliab Eng Syst Saf Dec. 2016;156:228–43.
[5] Mason P. A Bayesian analysis of component life expectancy and its implications on the inspection schedule. Reliab Eng Syst Saf May 2017;161:87–94.
[6] Malings C, Pozzi M, Klima K, Bou-Zeid E, Ramamurthy P, Bergés M. Surface heat assessment for developed environments: probabilistic urban temperature modeling. Comput Environ Urban Syst 2017;66:53–64.
[7] Malings, C. Pozzi, M. Klima, K. Bou-Zeid, E. Ramamurthy, P. and Bergés, M. “Surface heat assessment for developed environments: optimizing urban temperature monitoring,” 2017.
[8] Cressie NAC, Wikle CK. Statistics for spatio-temporal data. Hoboken, N.J.: Wiley; 2011.
[9] Raiffa H, Schlaifer R. Applied statistical decision theory. Cambridge, Massachusetts, USA: Harvard University Press; 1961.
[10] Malings C, Pozzi M. Value of information for spatially distributed systems: application to sensor placement. Reliab Eng Syst Saf Oct. 2016;154:219–33.
[11] Pozzi M, Der Kiureghian A. Computation of lifetime value of information for monitoring systems. In: Proceedings of the eighth international workshop on structural health monitoring. CaliforniaUSA: Stanford University; 2011.
[12] Straub D. Value of information analysis with structural reliability methods. Struct Saf 2014;49:75–86.
[13] Thöns S, Schneider R, Faber MH. Quantification of the value of structural health monitoring information for fatigue deteriorating structural systems. In: Proceedings of the twelfth international conference on applications of statistics and probability in civil engineering (ICASP12). Vancouver Canada; 2015.
[14] Malings C, Pozzi M. Conditional entropy and value of information metrics for optimal sensing in infrastructure systems. Struct Saf 2016;60:77–90.
[15] Memarzadeh M, Pozzi M. Integrated inspection scheduling and maintenance planning for infrastructure systems: integrated inspection scheduling and maintenance planning. Comput-Aided Civil Infrastruct Eng 2015;31(6):403–15.
[16] Memarzadeh M, Pozzi M. Value of information in sequential decision making: Component inspection, permanent monitoring and system-level scheduling. Reliab Eng Syst Saf Oct. 2016;154:137–51.
[17] Zejnilovic, S. Gomes, J. and Sinopoli, B. “Sequential observer selection for source localization,” 2015, pp. 1220–1224.
[18] Weerakkody S, Mo Y, Sinopoli B, Han D, Shi L. Multi-sensor scheduling for state estimation with event-based, stochastic triggers. IEEE Trans Autom Control Sep. 2016;61(9):2695–701.
[19] Koller D, Friedman N. Probabilistic graphical models: principles and techniques. Cambridge, MA: MIT Press; 2009.
[20] Malings C. and Pozzi, M. “Submodularity issues in value-of-information-based sensor placement,” in press.
[21] Memarzadeh M, Pozzi M, Zico Kolter J. Optimal planning and learning in uncertain environments for the management of wind farms. J Comput Civil Eng Sep. 2015;29(5):04014076.
[22] Bellman R. Dynamic programming. Princeton, NJ: Princeton Univ. Pr; 1984.
[23] Straub D, Faber MH. Computational aspects of risk-based inspection planning. Comput-Aided Civil Infrastruct Eng 2006;21:179–92.
[24] Krause A, Guestrin C. Optimal value of information in graphical models. J Artif Intell Res 2009;35:557–91.
[25] Oakley JE, Brennan A, Tappenden P, Chilcott J. Simulation sample sizes for Monte Carlo partial EVPI calculations. J Health Econ May 2010;29(3):468–77.
[26] Zitrou A, Bedford T, Daneshkhah A. Robustness of maintenance decisions: uncertainty modelling and value of information. Reliab Eng Syst Saf Dec. 2013;120:60–71.
[27] Daneshkhah A, Stocks NG, Jeffrey P. Probabilistic sensitivity analysis of optimised preventive maintenance strategies for deteriorating infrastructure assets. Reliab Eng Syst Saf Jul. 2017;163:33–45.
[28] Bensi M, Kiureghian AD, Straub D. Efficient Bayesian network modeling of systems. Reliab Eng Syst Saf Apr. 2013;112:200–13.
[29] Bartram G, Mahadevan S. Dynamic Bayesian networks for prognosis. In: Proceedings of the annual conference of the prognostics and health management society; 2013.
[30] Luque J, Straub D. Reliability analysis and updating of deteriorating systems with dynamic Bayesian networks. Struct Saf Sep. 2016;62:34–46.
[31] Memarzadeh M, Pozzi M, Kolter JZ. Hierarchical modeling of systems with similar components: A framework for adaptive monitoring and control. Reliab Eng Syst Saf Oct. 2016;154:137–51.
[32] Krause A. Optimizing sensing: theory and applications, 15213. Pittsburgh, PA: Carnegie Mellon University; 2008.
[33] Krause A, Guestrin C. Nonmyopic active learning of gaussian processes: an exploration–exploitation approach. In: Proceedings of the twenty-fourth international conference on machine learning. CorvallisOregonUSA; 2007.
[34] Golovin D, Krause A. Adaptive submodularity: theory and applications in active learning and stochastic optimization. J Artif Intell Res 2011;42:427–86.
[35] Papaioannou I, Betz W, Zwirglmaier K, Straub D. MCMC algorithms for subset simulation. Probab Eng Mech Jul. 2015;41:89–103.
[36] Rasmussen CE, Williams CKI. Gaussian processes for machine learning. Cambridge, Mass: MIT Press; 2006.
[37] Malings C, Pozzi M. Optimal sensing using value of information in spatio-temporal random fields. Presented at the 2016 working conference for the IFIP working group 7.5 on reliability and optimization of structural systems. PittsburghPAUSA; 2016.
[38] Malings C, Pozzi M. Optimal sensor placement and scheduling with value of information for spatio-temporal infrastructure system management. In: Proceedings of the twelfth international conference on structural safety and reliability, Vienna, Austria; 2017.
[39] Malings C, Pozzi M, Velibeyoglu I. Sensor placement optimization for structural health monitoring. In: Proceedings of the tenth international workshop on structural health monitoring; 2015.
[40] Robelin C-A, Madanat SM. History-dependent bridge deck maintenance and replacement optimization with Markov decision processes. J Infrastruct Syst Sep. 2007;13(3):195–201.
[41] Malings C, Pozzi M. Value of information analysis for typical bridge network topologies. In: Proceedings of the eighth international conference on bridge maintenance, safety, and management; 2016.