Reliability Engineering and System Safety 183 (2019) 93–103

Contents lists available at ScienceDirect

Reliability Engineering and System Safety

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

Submodularity issues in value-of-information-based sensor placement

C. Malings*, M. Pozzi

Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, United States

ARTICLE INFO

Keywords:
Value of Information
Entropy
Submodularity
Optimization

ABSTRACT

The value of information represents a rational metric for guiding the optimization of sensing efforts to support infrastructure system management under uncertainty. Unfortunately, this metric lacks the property of submodularity. Submodularity can be intuitively understood as a diminishing returns property, whereby the incremental benefit of a specific measurement is higher when the set of other available measures is smaller. Metrics which exhibit this property can be optimized using efficient greedy approaches, with certain guarantees on the near-optimality of the results. In this paper, we examine the issue of submodularity related to the optimization of sensor monitoring schemes using the value of information metric. We illustrate how greedy optimization approaches using value of information can lead to sub-optimal solutions for sensing in certain situations. We also examine how one potential heuristic approach involving a submodular surrogate metric (e.g. the conditional entropy) might be used to avoid some of these shortcomings.

1. Introduction

Infrastructure systems consist of many components, distributed across a spatial domain, which act together to fulfill a function. Over time, as the various internal (e.g. damage conditions) and external (e.g. applied loads) characteristics of the components change, the behavior of the system itself can change down to a lower capacity level or even a failure condition. Long-term management of the system, including maintenance or replacement of its components, is carried out to optimize its performance and functionality level. Various information can be collected during this process. In this paper, we use the term 'sensing' to generally refer to any information collection efforts within the system, including the installation of sensors or inspections by human or robotic agents. The collection of this information, in terms of schemes denoting the places and times at which information is gathered, is here referred to generally as "sensor placement and scheduling". This is performed strategically so as to maximize the impact of the collected information in supporting making better decisions about system management, considering also the costs associated with information collection. We give a qualitative overview of this problem of optimal sensor placement and scheduling in Section 2.

The value of information (VoI) metric is a decision-theoretic measure of the benefits of information in supporting decision-making in uncertain systems. Using knowledge about the various possible uncertain states of a system, the actions which can be taken to manage it, and the costs of various outcomes, the VoI quantifies how, in an expected sense, additional information can be used by a rational agent to improve the outcomes of decision-making [1]. VoI has been employed to support information gathering for civil infrastructure management, e.g. for quantifying the benefits of structural health monitoring [2,3] and of periodic inspections for deteriorating components [4,5]. While the VoI can be difficult to assess due to its computational complexity [6], specific assumptions on the structure of the decision-making problem can allow the metric to be efficiently computed for special cases, as discussed in the companion paper [7,8]. Although the present paper does not deal with the computation of the VoI metric, but rather with its optimization, a brief description of the metric is provided in Section 3, and an outline of the computational approaches developed in the companion paper is provided in Section 5.1.

In addition to the potential complexity of any evaluation of the VoI, its optimization related to sensor placement and scheduling can itself be a computationally daunting task. This is due to the exponential growth in the number of possible sensor placement and scheduling schemes as the number of available locations and the management time duration over which measures must be scheduled increase. In this paper, we examine the use of an efficient, approximate "greedy" optimization approach to solving the problem of sensor placement and scheduling. This approach iteratively adds or removes one measurement from a proposed scheme to maximize its net VoI, as if no further adjustments were possible. A similar greedy approach is applied for sensing based on conditional entropy and mutual information by [9–11] and based on VoI by [7,12]. However, as discussed in Section 4, there is no guarantee

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

https://doi.org/10.1016/j.ress.2018.11.010
Received 30 May 2017; Received in revised form 4 October 2018; Accepted 7 November 2018
Available online 10 November 2018
0951-8320/ © 2018 Elsevier Ltd. All rights reserved.

that such an approach will provide optimal or even provably near-optimal solutions under the VoI metric, due to the lack of submodularity for this metric. Submodularity is intuitively understood as a diminishing returns property, whereby the incremental benefit of a specific measurement is higher when the set of other available measures is smaller, and allows for theoretical guarantees on the near-optimality of greedy optimization approaches for metrics satisfying this property [13]. As illustrated in Section 5, a relevant example of the lack of submodularity of the VoI in infrastructure management is related to a common bias affecting a set of measures, in which case greedy approaches can perform quite poorly. For this reason, we propose and investigate the use of the submodular conditional entropy metric as a surrogate for the VoI in situations where the structure of a decision-making and sensing problem leads to poor performance by greedy approaches in maximizing VoI. Finally, some general conclusions on optimal sensing for the management of spatio-temporal systems are drawn in Section 6.

2. Problem statement

2.1. Optimizing sensing

The design of a sensing network or monitoring campaign can be viewed as a decision-making problem, where the decision variables are the sensor locations and/or the times at which measurements are to be collected. The set of chosen locations and times are referred to as the measurement scheme Y. Alternative schemes can be compared using an evaluation metric $\mathrm{M}(Y)$ which, taking the scheme Y as an input, assesses its overall benefit. The sensing scheme design process can then be formulated as a maximization of function $\mathrm{M}(Y)$. The difficulty of this maximization depends both on the difficulty of evaluating $\mathrm{M}(Y)$ for a given Y and the range of possible schemes Y which must be assessed.

Consider the placement of s sensors among r candidate locations. The number of distinct configurations is $\binom{r}{s}$, which grows rapidly; e.g., it is more than 17 trillion when $s = 10$ and $r = 100$. Hence, complete evaluation of all configurations is computationally prohibitive except for small systems, and more efficient algorithms need to be implemented.

2.2. Greedy approaches to combinatorial optimization

One classic approach to combinatorial optimization (i.e. the optimization of functions over a discrete, as opposed to a continuous, domain) is to make use of a "greedy" optimization algorithm. Two variations of this approach are considered in this paper. In a "forward greedy approach", sensors are placed one by one by solving a sequence of optimization problems involving a single decision variable each. Specifically, a sensor is first placed in the best location, considering it in isolation, according to metric M. Then, for $s - 1$ iterations, one additional sensor is placed in the best location, considering all previously placed sensors, again according to metric M. The corresponding search involves $rs - \frac{1}{2}s(s-1)$ evaluations of metric M, which is approximately rs evaluations for $s \ll r$. This is much less than what is needed by complete enumeration; e.g., only about 1000 evaluations are needed by the forward greedy approach when $s = 10$ and $r = 100$.

An alternative greedy scheme, the "reverse greedy approach" starts the search by placing sensors in all r candidate locations, and removes them one by one while maintaining the value of the metric as high as possible. This requires $r(r-s) - \frac{1}{2}(r-s-1)(r-s)$ metric evaluations. When $s \ll r$, this is approximately $\frac{1}{2}r^2$ evaluations, which would be much more than the approximately rs evaluations required by the forward approach (e.g., it is 5000 when $s = 10$ and $r = 100$). The number of evaluations needed by both approaches is equal when $s = \frac{1}{2}r$.

2.3. Submodularity

While the complexity of the greedy approaches scales well with the number of sensor placement options, these approaches are approximate, and so a question remains about their effectiveness: do we have any guarantee that the resulting choice of measurement scheme is satisfactory? This is strictly connected with the mathematical properties of function M. Although there are no general guarantees that the scheme selected by either greedy approach is truly optimal, for some metrics it is guaranteed that the greedily obtained solution provides a value of M close to the true optimum. The property of M that allows for this guarantee is called submodularity, and is intuitively related to the law of diminishing returns, as mentioned in Section 1. Formally, a set function $\mathrm{M}(\cdot)$ is submodular if, for every $Y_1 \subseteq Y_2 \subseteq \mathcal{Y}$ and every $y \in \mathcal{Y} \setminus Y_2$ (where $\mathcal{Y}$ denotes the powerset, or set of all possible $\mathcal{Y}$) it must hold true that [14]:

$$\mathrm{M}(Y_1 \cup \{y\}) - \mathrm{M}(Y_1) \geq \mathrm{M}(Y_2 \cup \{y\}) - \mathrm{M}(Y_2) \tag{1}$$

That is, the benefit of adding a new element y to the smaller set Y_1 is greater than the benefit of adding it to the larger set Y_2 (which includes Y_1). Note that this represents one of several equivalent definitions for submodularity.

Consider again the forward greedy approach: it starts placing one sensor in the optimal position when no other sensors are available. Could it be that this chosen location is not part of the optimal configuration when more sensors are used? This may happen, for example, because a more balanced use of many sensors can cover a wider area, including the chosen location, without the need to place a sensor exactly there. Intuitively, this is why a solution provided by the forward greedy approach might be suboptimal. However, in so far that the area has to also be covered with configurations using more sensors, the outcome of the greedy approach is only slightly worse than the optimal one. If this is the case, function M is submodular, and greedy approaches are effective.

However, sometimes a piece of information is relevant only when available in combination with others. Examples include a fragment of a larger code that needs to be completely known to work, or the latitude coordinate of a point on the Earth's surface (e.g., where a treasure is buried), when the corresponding longitude coordinate is also needed to guide the expedition. A metric which assigns a high value to a piece of information only in combination with others is not submodular. Generally, as mentioned in Section 1, VoI is not submodular. Examples of infrastructure monitoring applications when the VoI of a pair of measures lacks submodularity include the size and position of a crack, useful to decide repairs for maintaining the strength of a beam, or the concentration and chemical composition of a contaminant to devise an appropriate remediation strategy.

When the metric M is not submodular, the greedy approach may fail to discover configurations with high value. As a demonstration of this phenomenon, consider the two example measurement metrics reported in Table 1; these metrics quantify the "reward" obtained from placing sensors at various locations to measure a system. Here, $s = 2$ sensors can be placed among $r = 3$ possible locations, denoted A, B and C. Metrics M_I and M_{II} agree on the reward for placing a single sensor (3 units for A, 2 for either B or C), and in that of placing two sensors when one is at A (4 units for either A and B or A and C). However, they disagree on the reward for jointly measuring at B and C, and on that of

Table 1
examples of a submodular (M_I) and non-submodular (M_{II}) metric.

		measurement locations chosen						
		A	B	C	A,B	A,C	B,C	A,B,C
metric	M_I	3	2	2	4	4	3	5
	M_{II}	3	2	2	4	4	10	12

measuring all locations. In this example, both metrics assign a positive increment to the addition of a sensor to any set (e.g. both metrics assign an incremental benefit of 1 to adding a sensor at B to one at A). However, while metric M_I is submodular, metric M_{II} is not, since under this metric there is a greater benefit of adding a measure at location C to one at B (8 units) than that of using a measure at location C alone (2 units), and thus there are increasing rather than diminishing returns associated with this placement. The optimal configuration for two sensors according to M_I is either A, B or A, C; these sets can be obtained greedily by first selecting A, the single measure that maximizes the single sensor configuration, and then adding a second sensor. However, this greedy scheme will fail to capture the optimal configuration according to metric M_{II}, which is B, C. According to this second metric, while the benefit of monitoring just B or just C in isolation is low, there is a high benefit in monitoring them jointly, much higher than the sum of the individual benefits. The forward greedy approach will, again, start by selecting A, and so will prevent the discovery of the high value in monitoring the other two locations together. In this example, the reverse greedy approach would be able to identify the optimal configurations according to each metric; however, this better performance is not always guaranteed. Finally, we note that the optimization can also include, as an additional variable, the number of sensors s, if sensors are expensive and the target is maximization of the net benefit. Again referring to the example in Table 1, now suppose that the cost of placing 1 or 2 sensors is nil, while it is three units when three sensors are placed. Then, the optimal configurations will still be those described above, made up of two sensors.

Unfortunately, there is no hope of finding a generally effective search approach for problems of arbitrary complexity using arbitrary sensing metrics. As an example, consider a metric assigning a high value to the specific set A of k sensors, and no value to any other set, even if that set encompasses A. For this problem, no heuristic search method can be effective, as no sign of value can be detected outside A, unless this search can handle the simultaneous placement of k sensors. The only path that opens the possibility of finding an effective heuristic is based on the knowledge of the type of problem being addressed. In the following section, we examine the VoI metric for sensor placement, and then in Section 4, develop strategies for optimizing this metric while avoiding difficulties such as those discussed in this section.

3. Value of information

This section introduces the notational conventions of the paper and provides a brief overview of the VoI metric. Methods for efficient VoI evaluation have been developed in previous work [7,8,12]. Let F denote the set of random field variables which govern the performance of a system over space and time, with $f(\mathbf{x}, t)$ indicating the variable at spatial location $\mathbf{x}$ and time t. Multiple co-located random fields, e.g. temperature and humidity at the same spatio-temporal coordinate, can be modeled by augmenting the spatial location vector to include a field index. These variables are modeled via an appropriate spatio-temporal random field model p_F describing the joint prior distribution of these quantities [15]. Let vector $\mathbf{f}$ include all of the random field variables affecting the system at a finite number of spatial locations and finite number of discrete timesteps over which the performance of the system is to be modeled.

A measurement scheme Y defines a set of locations and times for measuring these random field variables to observe the system and control it. This scheme is selected as a subset of $\mathcal{Y}$, the collection of all possible measurement locations and times. The cost of implementing this scheme is given by the sensing cost function $C(Y)$. The realized collected measures are denoted by observation vector $\mathbf{y}$. Appropriate Bayesian methods allow this data to define a posterior distribution $p_{F|\mathbf{y}}$ for the random field variables conditional to these measurements. Based on the prior or posterior model (as appropriate) a managing agent must make decisions about how to control the system. Let these decisions be encoded into a vector $\mathbf{a}$ of actions (potentially including empty or 'do nothing' actions), which are selected from a set $\mathcal{A}$ of possible options. The choice of management actions, together with the values of the random field variables, uniquely define a loss, i.e. a penalty (or negative benefit) incurred by the managing agent for any specific outcome of the decision-making process. This is quantified via loss function $L(\mathbf{f},\mathbf{a})$. Typically, the loss is quantified in monetary terms, although any scalar utility measure can be applied.

For any agent without any access to additional information, the best open-loop policy is to select a set of actions which minimizes the expected value of the loss function over the possible states of random variables, obtaining:

$$\mathbb{E}L(\varnothing) = \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_F L(\mathbf{f}, \mathbf{a}) \tag{2}$$

where $\mathbb{E}_F$ denotes the statistical expectation with respect to p_F.

If, on the other hand, the agent has access to additional information, this information should be used in a closed-loop policy, updating the prior model and revising the set of selected actions accordingly. Before observations are made, their impact can be predicted in an expected sense, by taking the expectation (over possible observation outcomes) of the minimum (over possible management actions) of the expected loss given the observations. This is denoted as the posterior expected loss, and calculated as:

$$\mathbb{E}L(Y) = \mathbb{E}_Y \min_{\mathbf{a}\in\mathcal{A}} \mathbb{E}_{F|\mathbf{y}} L(\mathbf{f}, \mathbf{a}) \tag{3}$$

where $\mathbb{E}_{F|\mathbf{y}}$ denotes the statistical expectation with respect to $p_{F|\mathbf{y}}$. Note that this formulation assumes that all observations are made before any management actions are taken. If this is not the case (e.g. if measurements of the system at earlier times are available to support the selection of later management actions) the posterior expected loss must be computed recursively, as illustrated in [8].

VoI is defined as the difference of the prior and posterior expected losses for managing the system:

$$\mathrm{VoI}(Y) = \mathbb{E}L(\varnothing) - \mathbb{E}L(Y) \tag{4}$$

In other words, the VoI quantifies how much the loss will decrease, in an expected sense, if decisions are made after considering additional information collected by the measurement scheme Y [1]. VoI is also sometimes referred to as the expected net gain by sample information (ENGS) or expected value of sample information (EVSI). Details of the evaluation of the VoI in general are provided in the companion paper [8].

4. Strategies for sensor placement and scheduling

An optimal sensing scheme according to the VoI metric is a scheme Y, selected from the set $\mathcal{Y}$ of all possible schemes, which maximizes the net VoI, i.e. the VoI minus the cost of collecting measurements. This is expressed as:

$$Y^* = \mathrm{argmax}_{Y\subseteq\mathcal{Y}}\ \mathrm{VoI}(Y) - C(Y) \text{ subject to } C(Y) \le b \tag{5}$$

where b represents a fixed sensing budget. This represents a choice of metric $M(Y) = \mathrm{VoI}(Y) - C(Y)$, following the discussion of Section 2.1. Problems of optimal sensor placement and/or scheduling can be addressed through appropriate definition of the sensing cost function to capture constraints. For example, a sensor placement problem can be addressed by assigning negligible additional costs to multiple measurements selected for the same location across different timesteps, while simultaneous placement and scheduling problems should assign appropriate costs based on the rates at which measurements are made, even for common spatial locations.

4.1. Forward and reverse greedy optimization

As discussed in Section 2, the exact solution of Eq. (5) is in general

Algorithm 1
Pseudo-code for the forward greedy algorithm.

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

intractable, and more efficient approximate solution approaches must be employed; in this paper, two greedy different optimization algorithms are illustrated for this purpose. Both algorithms reduce the computational demand of the optimization from $\binom{r}{s}$ evaluations of the VoI to less than rs evaluations in an optimal set of size s is chosen out of r candidate sensing locations. The forward greedy optimization algorithm iteratively builds the scheme Y^* by, at each step, selecting the individual measurement from the set of unselected measurements in $\mathcal{Y}$ which, when combined with previously selected measurements, maximizes the value of the optimization objective. Algorithm 1 presents pseudo-code for this forward greedy algorithm. Note that in the code this algorithm begins with an empty set and proceeds to build increasingly larger proposed optimal sets until the entire set of possible measurements is exhausted. Then, all proposed optimal sets are compared and the set which maximizes the objective is chosen. This is necessary because the non-submodular objective is not necessarily non-decreasing as the number of measurements increase (this will be discussed in more detail later). However, in practice, the algorithm can typically be halted prior to the exhaustion of the full set of potential measurements, i.e. when the given sensing budget has been exhausted.

An alternative to the forward greedy optimization approach is the reverse greedy approach, as presented, e.g., in [16]. This algorithm starts with the set of all possible measurements $\mathcal{Y}$ and iteratively removes measures while keeping the objective function as high as possible. Algorithm 2 provides the corresponding pseudo-code. Note that this algorithm proceeds to generate proposed optimal sets of smaller sizes until the null set is selected, and then trims away any sets which exceed the sensing budget before comparing the remaining sets to determine the optimal solution. In this case, there is no opportunity to stop the algorithm when the budget has been exhausted, as the algorithm must proceed from the set of all possible measures. Computationally, therefore, the reverse greedy approach is less efficient than the forward greedy approach, especially considering that in many cases the size of the optimal set is much smaller than that of the set of all possible measures, i.e. $|Y^*| \ll |\mathcal{Y}|$.

Finally, note that, as presented here, both the forward and reverse

Algorithm 2
Pseudo-code for the reverse greedy algorithm.

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

greedy approaches rely on the discretization of the spatial and temporal domain of measurements, such that the set $\mathcal{Y}$ is of a finite size. In general, continuous sensing domains can also be handled via a forward greedy search approach. For example, the positioning in space and scheduling in time of a measurement can be performed via a gradient-based approach using the VoI metric. This procedure can be repeated to add additional measurement locations and times in a greedy manner, and continued until an appropriate stopping point, e.g. budget exhaustion. The reverse greedy approach is no longer applicable, however, since it is impossible to define any finite set of all possible measurements in a continuous domain.

4.2. Submodularity and a submodular surrogate metric

For metric functions which have the property of submodularity, greedy optimization approaches are guaranteed to achieve at least $1 - \frac{1}{e} \cong 63\%$ of the metric value of the true optimal solution obtained via an exhaustive search [13]. Unfortunately, there are no general guarantees on optimality of results to Eq. (5) obtained via either of the greedy approaches discussed above due to the fact that the VoI metric does not in general satisfy the property of submodularity. This can be demonstrated through an intuitive example. Suppose an agent prefers to catch a train instead of using their own car for a long and urgent trip, but cannot risk arriving late to their destination. Two pieces of information are available to support the decision of whether or not to take the train: checking a timetable to determine when the train will leave the station and checking local traffic conditions to determine when is the earliest one might reach the station via taxi. Either of these pieces of information, in isolation, cannot help the agent much, as they cannot risk missing the train. However, the combination of these data will allow them to take an informed decision. Therefore, while the VoI of each piece of information, in isolation, is negligible, the VoI of the combination is high. Through this example, as well as those of Section 2.3, it can be seen that VoI may generally violate the property of submodularity, and therefore a greedily optimized information gathering plan based on this metric may perform arbitrarily badly as compared to the best possible plan. Generally, a greedy optimization plan will perform poorly when several pieces of data have a high VoI when combined, but have a much smaller or negligible VoI individually, as discussed in Section 2.

Past work has developed efficient approaches to optimize measurement selection in chain graphical models, such as hidden Markov models, but generalization of these techniques to other model structure such as polytrees (i.e. a directed acyclic graph with a tree structure, a subset of Bayesian Network PGMs) has been found to be impractical in general [17,18]. Prior work on optimal sensor placement has found that the greedy approach can provide suitable solutions to practical sensor placement problems using the VoI metric even without guarantees on the near–optimality of results e.g. [19]. Therefore, in this work, greedy optimization approaches are employed to allow for efficient selection of sensing schemes, and are also compared to alternative non-greedy approaches. The examples of Section 5 illustrate the application of these methods, as well as demonstrate possible shortcomings of greedy approaches in optimizing the non-submodular VoI metric.

Considering the issues regarding the lack of submodularity of the VoI metric, and its potential effects on the quality of greedily optimized sensing schemes, the identification of a submodular sensing metric that can serve as a surrogate for VoI is of great interest. In this paper we make use of the conditional entropy metric as the submodular surrogate. Conditional entropy is an information theoretic measurement of the uncertainty in a set of random variables given observations of a related set of variables, and is known to be submodular in the case that observations Y are conditionally independent given the random field variables F [9,20,21]. We denote by $\mathrm{H}(F)$ the entropy of the random field variables affecting a system, and by $\mathrm{H}(F|Y)$ the conditional entropy of these variables given measurements taken according to scheme Y.

This entropy is evaluated as:

$$\mathrm{H}(F|Y) = -\mathbb{E}_{F,Y}\log(\mathrm{p}_{F|\mathrm{y}}) \tag{6}$$

Note, however, that at any given time only observations made prior to that time can be utilized to reduce uncertainty in the system. Therefore, for a given timestep, the uncertainty reductions resulting from planned future observations should not be counted towards the overall uncertainty reduction at that timestep. Considering that decisions are made sequentially at each timestep, it is the cumulative uncertainty of the random field variables across these timesteps, conditioned on sets of past measured data, which is better related to the uncertainties faced by agents managing these systems. For this reason, our proposed surrogate metric for sensor placement is:

$$\mathrm{M}_{\mathrm{Ent}}(Y) = m\mathrm{H}(F) - \sum_{j=1}^{m}\mathrm{H}(F|Y_{\rightarrow j}) \tag{7}$$

where $Y_{\rightarrow j}$ denotes the subset of measurements in Y collected before timestep j, and m is the total number of timesteps. That is, the metric is the difference between the sum of prior entropies and the sum of the conditional entropies of the random field variables over the management time duration. Note, however, that the sum of prior entropies is constant with respect to Y, and therefore plays no role in the optimization.

In this paper, instead of than applying the conditional entropy metric directly to determine an optimal sensing scheme (as this can itself produce suboptimal results according to the VoI metric [7]), it is combined with the VoI metric in a modified optimization heuristic based on the forward greedy algorithm, as a way to "break" the algorithm out of suboptimal solution paths. The motivation for and general operation of this heuristic methods is as follows. During successive iterations of the forward greedy algorithm, the amount of additional benefit provided by each new selected measurement will occasionally be low. During such times, it may be beneficial in the long-term to select measurements which, while suboptimal by the VoI, may be combined with other measurements in later iterations to provide a large VoI. Such a scenario is possible due to the non-submodularity of the VoI metric. Rather than selecting these next measurements randomly, the conditional entropy metric can be used to rationally guide the selection of new measurements which will most reduce uncertainty in the system. While this uncertainty reduction will not immediately lead to better decision-making, in the long-term the reduction of uncertainty can increase the value provided by other measures and lead to a better overall measurement set. This heuristic approach is presented formally in Algorithm 3, where the surrogate conditional entropy metric is used to select the next measurement to be added whenever the measurement suggested by the VoI metric will not lead to an increase in the net VoI (the VoI minus sensing cost) compared to the previously selected measurement set. In this way, the sensing cost is used to establish the threshold for unacceptably low VoI increases. Note that this choice of threshold may not always be appropriate, e.g. in cases where the costs of certain measurements are negligible, such that a large number of relatively useless measures might be selected. Use of an alternative virtual cost in these cases, reflecting a small additional cost imposed with every measurement, would be a way to avoid this difficulty. Section 5 presents examples of the use of this heuristic method which help to illustrate how it can help to "break" greedy forward searches based on VoI alone out of suboptimal paths.

5. Application examples

This section includes examples illustrating the application of the optimal sensor placement methods discussed in this paper. Section 5.1 discusses the efficient evaluation of VoI in spatio-temporal systems, referring to the results of [8]. Section 5.2 provides an example comparing forward and reverse greedy optimization, and illustrating the shortcomings of greedy optimization approaches for non-submodular metrics like VoI. Section 5.3 presents an in-depth example application of the optimization methods of this paper, as well as assesses the performance of the surrogate conditional entropy metric for sensor placement optimization. Both examples of this section make use of Gaussian process models; definition and notation for these models are described in Appendix A.

Algorithm 3
Pseudo-code for the forward greedy algorithm, including submodular surrogate objective function.

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

5.1. Efficient value of information evaluation

In general, computation of the VoI is complicated by the fact that the dimensionality of the problem space increases exponentially in the number of possible joint system states, the number of potential management actions for the system, and the number of observational outcomes. However, under specific assumptions about the structure of the loss function, the computational demand can be reduced significantly. In particular, the loss function may be decomposable in space (i.e., the function decomposes across different parts or components of the system), in time (i.e., the function decomposes over discretized timesteps of the system management time duration), or in both space and time. We denote such a decomposable loss function as follows:

$$\mathrm{L}(\mathbf{f},\,\mathbf{a}) = \sum_{j=1}^{m}\gamma_j\sum_{i=1}^{n}\mathrm{L}_{i,j}(\mathbf{f}_{i,j},\,\mathbf{a}_{i,j}) \tag{8}$$

where $\mathbf{f}_{i,j}$ refers to a subset of $\mathbf{f}$ corresponding to the random variables affecting a single location or component of the system indexed by $i \in \{1, \ldots n\}$ at timestep $j \in \{1, \ldots, m\}$, $\mathbf{a}_{i,j}$ refers to the subset of $\mathbf{a}$ corresponding to actions taken to manage the system at this location and time, and γ_j is a positive scalar factor used to discount losses for timesteps to a common reference time (in case such discounting is appropriate to the system management timescale being considered). Under this assumption, the loss function can be expressed as the (possibly discounted) sum of local loss functions corresponding to different times and locations within the system.

Furthermore, under this assumption, the actions taken to manage the system have only local effects. That is, the choice of actions corresponding to one location cannot influence the loss function associated with another location, and the choice of actions at a specific time cannot affect the performance of the system at later times. Such a system is termed as an uncontrolled system, because the actions taken to manage the system do not control its evolution in future timesteps. While such an assumption can be restrictive, especially in systems

managed over a long time duration, the assumptions of decomposability in the loss function and locality of action effects allow the evaluation of the VoI to be performed for each location and timestep separately, and the results summed. Since each individual evaluation involves only a small subset of the actions and variable states, the dimensionality of these subsets is greatly reduced. This dimensionality will grow linearly in the number of system components or locations n and in the management time duration length m. The evaluation of the VoI can therefore be performed much more efficiently. The details of VoI computation under these assumptions (and more generally) are provided in the companion paper [8].

For defining the submodular surrogate metric of Section 4.2 in systems with decomposable loss functions, $\sum_{i=1}^{n} \mathrm{H}(F_{i,j}|Y_{\rightarrow j})$ is substituted for $\mathrm{H}(F|Y_{\rightarrow j})$ in the definition of Eq. (7). This is done for two main reasons. First, the losses associated with each timestep and location are only functions of the random variables and actions for that timestep and location, and actions are selected based only on these random variables, so only the uncertainties in these random variables impact the decision-making process at each timestep and location. Second, this formulation reduces the computational demands of the metric evaluation by reducing the number of variables whose joint entropy must be evaluated.

Finally, it is important to note that, although it is convenient to make use of decomposable systems when evaluating VoI, it is by no means necessary to do so. The algorithms and heuristics presented in this paper are completely agnostic to the manner in which the VoI is evaluated, and will be equally applicable to other systems in which the assumptions stated above do not apply.

5.2. Comparison of forward, reverse, and heuristic-based greedy optimization

This example compares forward and reverse greedy optimization of the VoI metric, as well as the proposed heuristic approach of Section 4.2, in an intuitive application. The structure of this problem is motivated by an example of differential settlement between columns which has been investigated in previous work [22]. Consider a system with two components whose states change over time, which can be monitored by precise but biased sensor systems. This example will demonstrate how correlated noise (i.e., bias) in measurements can lead the VoI to exhibit non-submodular behavior, and how this behavior can cause suboptimal performance in greedy sensor selection. Specifically, in this example, useful information will derive from pairs of measurements (one measure of the state and one to remove the bias) on either component; the forward greedy approach, which is unable to assess the value in such pairs, will thus pursue a suboptimal sensing strategy; it will in fact focus on measuring one component to the exclusion of the other.

The behavior of the system over $m = 10$ timesteps is described with $n = 2$ random variables, one associated with either component. These random variables are described via a Gaussian process model with zero mean, i.e. $\mu(\mathbf{x}, t) = 0$, and with covariance function:

$$\mathrm{k}(\mathbf{x}, t, \mathbf{x}', t') = \begin{cases} \sigma(t)\sigma(t')\exp\left(-\dfrac{|t-t'|^2}{2\lambda^2}\right) & \text{if} \mathbf{x} = \mathbf{x}' \\ 0 & \text{if} \mathbf{x} \neq \mathbf{x}' \end{cases} \tag{9}$$

where $\sigma(t) = 0.5 - 0.5\exp(-t)$ and $\lambda = 5$. This models a system where the performance of the components are spatially uncorrelated but the performance of each component is temporally correlated, with uncertainty in the performance (i.e. the variance of the random variables) increasing over time. In this example, space, time, and random field values are dimensionless quantities; in general, such values can be described using any appropriate system of units for the problem, e.g., meters, seconds, kilograms.

It is assumed that measurements of these random variables may be taken at any timestep, but these measurements are associated with a large, highly correlated error which is common to all measures taken at the same location. For example, this may represent measurements which are taken in relation to a common initial reference point, the location of which is not known beforehand. This error is modeled as Gaussian noise as in Eq. (18) with $\boldsymbol{\mu}_\epsilon = \mathbf{0}$ and with the entry in the kth row and l^{th} column of the covariance matrix being:

$$\boldsymbol{\Sigma}_{\epsilon(k,l)} = \sigma_\epsilon^2 \, \mathbb{I}(k = l) + \sigma_{\text{bias}}^2 \, \mathbb{I}(\mathbf{x}_{Y_k} = \mathbf{x}_{Y_l}) \tag{10}$$

where $\mathbb{I}(\cdot)$ is an indicator function, taking on value 1 if its argument is true and 0 otherwise, and $\mathbf{x}_{Y_k}$ indicates the location associated with the kth measure, i.e. the index of the component being measured. This is used to model a situation in which measures associated with each component are precise, with noise level $\sigma_\epsilon = 10^{-4}$, but are subject to a large unknown bias, with standard deviation $\sigma_{\text{bias}} = 100$ (error with a high correlation across time is equivalent to a systematic bias in the measurement, as discussed in Appendix A).

Management of this system is defined using a temporally decomposable loss function and an uncontrolled system. The decomposable loss of Eq. (8) is used with:

$$\mathrm{L}_{i,j}(\mathbf{f}_{i,j}, \mathbf{a}_{i,j}) = \begin{cases} 100 & \text{if} |\mathbf{f}_{i,j}| > \delta_{\max} \text{and} \mathbf{a}_{i,j} = 0 \\ 20 & \text{if} \mathbf{a}_{i,j} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

where $\delta_{\max} = 0.5$. This reflects a problem in which separate binary actions are taken to manage each component. If the absolute value of the random variable associated with a component exceeds a certain threshold, a heavy loss is incurred, while if a management action is implemented for the component, a reduced loss is incurred regardless of the random variable state. For example, this might represent a temporary corrective action applied to a deteriorating component. The action prevents the full consequences of component failure from being realized, but must be renewed periodically. This contrasts with a corrective action like permanently replacing a component, which would violate the assumptions necessary for efficient computation of the VoI using the methods of Section 5.1. No discounting is assumed for the problem, i.e. $\gamma_j = 1 \;\; \forall \; j \in \{1, ..., m\}$, and the costs of measures are ignored, i.e. $\mathrm{C}(Y) = 0$. All losses and costs are dimensionless for this problem; in general, these may be quantified in monetary terms.

Fig. 1 demonstrates the results of applying both forward and reverse greedy optimization approaches to the problem of sensor placement and scheduling for this system under the VoI metric. Fig. 1a shows the order in which measures associated with both components are chosen by the forward (numbers above the lines) and reverse (numbers below the lines) greedy optimization approaches. For example, the optimal set consisting of ten measures selected by the forward greedy approach consists of the ten measures on component 1 for all ten timesteps, while the set of the same size selected by the reverse approach consists of measures on both components at timesteps $t = 0$, $t = 1$, $t = 2$, $t = 5$, and $t = 7$. Fig. 1b indicates the VoI for sets with different numbers of measures selected via either approach, as well as via the proposed heuristic approach of Algorithm 3. For example the set with ten measures selected by the forward greedy approach has a VoI of about 120, while the set of the same size chosen by the reverse approach has a VoI of about 240.

Examining the sequence of measures selected by the forward greedy approach, one can appreciate the pitfall to which this approach has succumbed when applied to this non-submodular metric. Initially, because of the high unknown bias of all measures, there is negligible value associated with any individual observation. As a result, a measure associated with the first component at the first timestep is arbitrarily selected. However, after this measure is collected, and because it is known from the prior model that the random variable values associated with this timestep are close to 0, the bias associated with all subsequent measures of the variable associated with this component can be determined and eliminated. This causes the value provided by subsequent

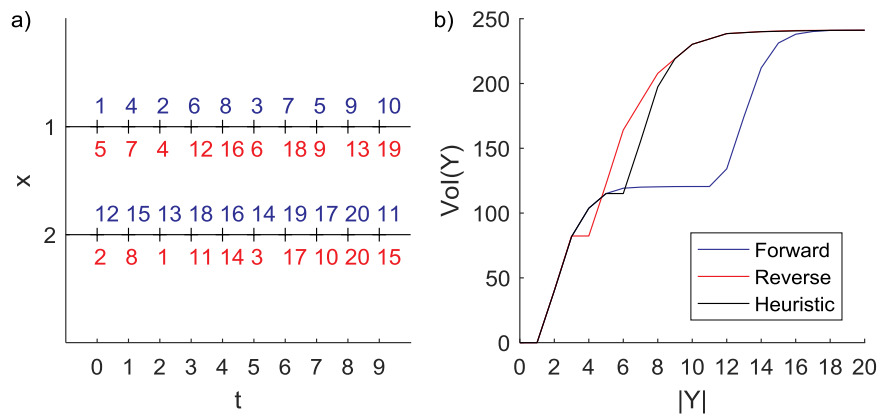

Fig. 1. a) Order of greedy selection of measures by forward (above lines) and reverse (below lines) greedy optimization; b) VoI versus number of sensors for forward and reverse greedy optimization, as well as the proposed heuristic approach.

measures of this component to increase. For the next several iterations of the algorithm, measures associated with the first component are selected until all such measures are exhausted. However, the value provided by later measures is diminished because of the high redundancy in the collected information. When the algorithm considers measurements on the second component, the VoI for these measurements is negligible; the algorithm is unable to consider that taking such a measurement would provide a reference point against which future measurements could be compared, thus increasing their value. Instead, the minimal VoI provided by additional measures on the first component outweighs the negligible VoI provided by any single measure on the second component, until all measures on the first component have been exhausted and the algorithm has no choice but to move to the second component. Note that, when this occurs, a measure at the last timestep is arbitrarily selected first, but the next selected measure is at the first step. Again, the reason for this is that a pair of measures on a common component (one at the first timestep to identify and correct the bias and one later on to support management decision-making) has higher value than either measure individually, a consequence of the lack of submodularity of the VoI metric.

The reverse greedy optimization approach largely avoids this pitfall in this example. By starting the algorithm with the full set of measures, the advantage of measures on both components at the first timestep to correct for the bias in subsequent measures is captured from the start. These initial measures are thus preserved, while redundant measures are eliminated. Only when the total number of sensors is relatively low is one component neglected (from the monitoring point of view), whereas in the forward approach the second component is neglected until half of all possible measurements are selected. One can imagine this problem being exacerbated by extending the time duration further, allowing the forward approach to linger even longer on the first component.

Although the reverse greedy optimization approach outperforms the forward approach in this example, it is not necessarily superior in other contexts. As has been mentioned, the need to begin with the full measurement set and eliminate measures increases the practical computational difficulty of the reverse approach. In this problem, due to the relatively small number of possible measures, this shortcoming did not play a significant role. Furthermore, note that the forward approach actually outperforms the reverse approach for a set of four measures, and thus the reverse approach does not dominate the forward, even in this example. In general, there is no guarantee that either algorithm will outperform the other.

The heuristic optimization approach based on the submodular surrogate (conditional entropy) metric is also illustrated for this problem. This approach follows the forward greedy optimization exactly until after the fifth measure has been chosen for the first component. At this point, the additional gain of VoI from adding a sixth measure for this component, as suggested by the forward greedy algorithm, falls below a predefined threshold (in this case 10 units) and the surrogate metric is invoked. The surrogate selects a measure on the second component at time $t = 0$, as this would have the least in common with all previously selected measures, and thus lead to the largest conditional entropy reduction for the system. With this new measure added, the algorithm now switches back to VoI. At this point, because an initial measurement on the second component has now been included, the values of subsequent measures for this component increase, and the VoI continues selecting measures there, quickly bringing the total VoI up to the levels achieved by the reverse greedy approach. After the twelfth sensor is placed, VoI gain again drops below the defined threshold and conditional entropy is used to guide the selection of the last group of measures. However, because VoI is already close to its maximum achievable value (attained by including all 20 possible measurements), there is little difference between the VoI achieved by this heuristic and that of either the forward or reverse greedy approaches on these final few measurements.

This example has illustrated how the lack of submodularity of the VoI metric can have a significant effect on the performance of a greedy optimization approach in an intuitive example application involving biased measurements. Although the reverse approach outperforms the forward approach in this example, avoiding the main pitfall to which the forward approach was susceptible, it is still not dominant, and may generalize poorly to problems with larger sets of potential measures. On the one hand, an intuitive understanding of the problem structure and its effect on the optimization algorithm can help guide the selection of a more appropriate optimization approach, e.g. by recognizing the need to correct for biased measures and performing this correction manually before selecting subsequent measures via a greedy approach. On the other hand, in larger and/or more complicated problems, such an intuitive understanding may be elusive, and recognizing and avoiding pitfalls to which the approaches may fall victim is more difficult. For this reason, the proposed heuristic method, which detects and attempts to “break free” of suboptimal solutions, may provide a reasonable and efficient solution, as it does in this example.

Finally, the heuristic approach is compared with two other alternative approaches to sensor placement optimization for this problem. For this comparison, a uniform cost of 10 units is assigned to each measure, such that the overall optimal result (i.e. that which maximizes the VoI minus the sensing cost) will not consist of all measures. First, the conditional entropy metric is used directly to perform the

optimization (also utilizing the forward greedy algorithm). Note that conditional entropy is evaluated for all variables involved in the problem, including the unknown bias of the measurements. Second, a commonly used "off-the-shelf" approach to solving complicated combinatorial optimization problems, a genetic algorithm, is used. Genetic algorithms attempt to identify an optimal solution by testing many randomly generated potential solutions and combining the best of these in an attempt to generate even better solutions for future iterations. However, there are again no guarantees that such an approach will achieve an optimal solution [9]. Genetic algorithms have been used for combinatorial optimization and sensor placement, e.g. by [23], where they showed promising results. For this problem, a genetic algorithm solver implemented in MATLAB (as the "`ga`" function) is used. The goal given to the algorithm is to maximize the net VoI (i.e. VoI minus the sensing cost) of a set of measurements; the algorithm may select any combination from among the 20 possible measurement locations and times depicted in Fig. 1a. The "`ga`" function by default uses a population size of 200, evolves the population using randomized crossovers (with 80% of the population generated by crossovers) and randomized mutations, and maintains the best-performing 5% of the population between generations; these default options were maintained in this example, and the reader is referred to the MATLAB documentation for further details (https://www.mathworks.com/help/gads/ga.html). For the genetic algorithm, two separate runs are performed. First, the algorithm is run several times, with each run constrained to select a measurement set of a specified size and to complete within the runtime of the heuristic approach for a set of the same size; this allows an equitable direct comparison of these two methods. Second, the algorithm is run in an unconstrained (both with respect to number of measures and runtime) manner to select a single optimized measurement set of arbitrary size. This is done since a limited run-time would not necessarily allow the algorithm to achieve convergence. In this unconstrained case, the default stopping criteria of the "`ga`" function are used instead, i.e. that the algorithm stops when there has been no significant (relative change in objective function values below about 10^{-5} percent) improvement in the objective function for 50 generations, or when the total number of generations exceeds 100 times the number of variables to be optimized (in this case, with the variables to be optimized representing whether or not to include each possible measurement location and time in the selected set, the maximum number of generations is 2000).

Fig. 2 illustrates the results of these approaches. Fig. 2a depicts the sets of measures belonging to the optimal sets chosen by each approach; measures belonging to the optimal set selected by the heuristic approach are depicted as black dots, those selected by the pure entropy metric (i.e. using Algorithm 1 with $\mathrm{M}_{\mathrm{Ent}}(Y)$ instead of $\mathrm{VoI}(Y) - \mathrm{C}(Y)$ as the objective function) are purple circles, and those selected by the genetic algorithm (for its unconstrained run) are green triangles. Fig. 2b plots VoI as a function of number of sensors in the selected set, and the set which maximizes the net VoI is indicated by the highest peak. Note that, while the heuristic and genetic algorithm approaches directly optimize the VoI metric, the conditional entropy approach does not. Instead, greedy forward optimization is performed based on pure entropy alone, and the VoI associated with each of the proposed optimal sets of different sizes is evaluated. For the genetic algorithm, the overall optimal solution resulting from the unconstrained run is indicated by the green star, while the green line indicates results of the constrained runs for different measurement set sizes.

In this simple example all three approaches yield similar optimal results, with the heuristic method and the unconstrained genetic algorithm achieving different solutions with nearly identical net VoI. For the conditional entropy metric, in this case, the chosen optimal set consists of all measures on both components before time $t = 5$. This neglects later measures, which are important for the management of the system at later times, and so the overall VoI is lower. When subjected to runtime constraints, the genetic algorithm does not perform as well overall, but often gives comparable results to the heuristic method; however, in a problem with a larger solution space to be searched, this difference might be more significant.

5.3. Application of the heuristic approach in a larger problem

In this second example, we consider the displacement of a system distributed over space, such as the movement of a track or conduit from its ideal path; an illustration of this is provided in Fig. 3. We consider that the correct positioning of this system vertically and horizontally along its length is essential to its proper functioning. In this problem, certain measurements are again highly biased, and must be paired with appropriate "calibrations" to provide significant value. By assigning costs to these various measurements, we are able to illustrate how the forward greedy approach can overspend on less informative single measurements while missing the potential for a greater return on investment from valuable pairs of measures. We also illustrate how the surrogate metric can "break" the forward greedy algorithm out of this suboptimal path. The dimensionality of this problem, in terms of the number of possible measurements considered, is too large for a feasible application of the reverse greedy approach (the reverse approach would take about ten times longer than the forward approach in this case); therefore, only forward greedy approaches are feasible for this problem.

We discretize the system into $n = 120$ segments over a 20m linear

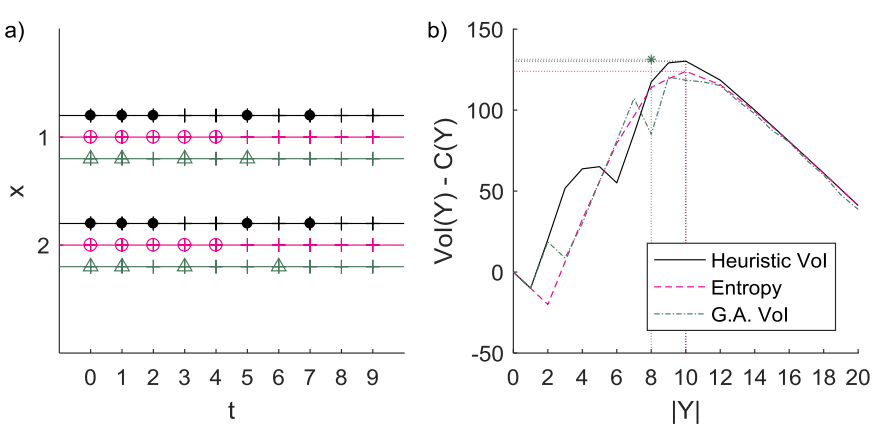

Fig. 2. a) Optimal measurement sets selected by the proposed heuristic approach (black dots), along with those of the conditional entropy metric alone (purple circle) and those selected by a genetic algorithm using VoI as the optimization metric (green triangle), with results from each algorithm having a slight vertical offset to improve clarity; b) net VoI (VoI minus cost) versus number of sensors for the proposed heuristic approach (black), as compared to conditional entropy (purple) and genetic algorithm (green) solutions. Locations of peaks are denoted by dotted lines. For the genetic algorithm, optimized sets of various sizes resulting from constrained runs are represented by the green line, while the overall solution obtained for a separate unconstrained run (giving the results depicted to the left) is indicated by the green star. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

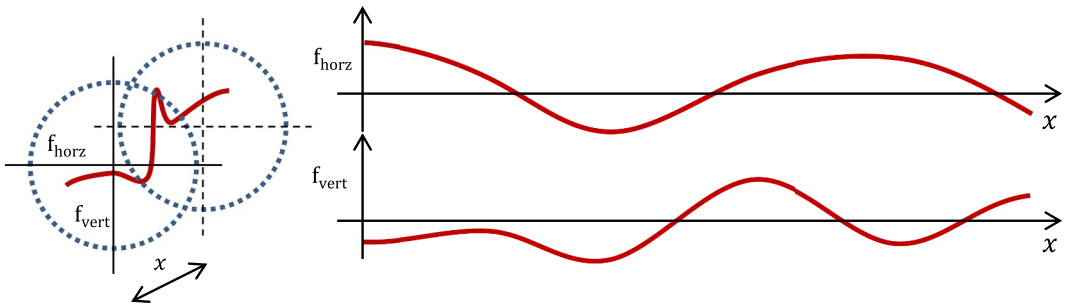

Fig. 3. Illustration of the example. The horizontal (f_{horz}) and vertical (f_{vert}) displacements of the system from its ideal position vary along its spatial extent (x).

domain, and model the vertical and horizontal displacements of these segments from their initial positions using a Gaussian process random field model. This model has a zero mean in both axes and a spatial covariance function as in both the vertical and horizontal directions as follows (with the formula or the horizontal direction being the same as for the vertical direction, shown here):

$$k_{vert}(x, x') = \sigma_X^2 \exp\left[-\frac{(x-x')^2}{2\lambda_X^2}\right] \tag{12}$$

where $\sigma_X = 1\text{cm}$ and $\lambda_X = 2\text{m}$. This function models the spatial correlation of displacements with a square exponential covariance function. The displacements in the vertical and horizontal directions are assumed to be independent.

In this problem, only a single management timestep is considered, i.e. $m = 1$, modeling a sensor placement problem only. For the management of the system, a penalty is incurred if the vertical or horizontal displacement exceeds 2cm, with this penalty being additive across the system. Mitigation actions are possible to correct vertical or horizontal displacements. The loss function encoding this problem is:

$$L_{i,j}(\mathbf{f}_{i,j}, \mathbf{a}_{i,j}) = L_{i,j,vert}(\mathbf{f}_{i,j,vert}, \mathbf{a}_{i,j,vert}) + L_{i,j,horz}(\mathbf{f}_{i,j,horz}, \mathbf{a}_{i,j,horz}) \tag{13}$$

where $\mathbf{f}_{i,j,vert}$ indicates the vertical displacement at discrete location i and timestep j, and $\mathbf{a}_{i,j,vert}$ indicates the corresponding management action. The loss functions regarding vertical and horizontal displacement have a similar form. For example, for vertical displacement:

$$L_{i,j,vert}(\mathbf{f}_{i,j,vert}, \mathbf{a}_{i,j,vert}) = \begin{cases} 25 & \text{if} |\mathbf{f}_{i,j,vert}| > \delta_{max} \text{ and } \mathbf{a}_{i,j,vert} = 0 \\ 5 & \text{if } \mathbf{a}_{i,j,vert} = 1 \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

where $\delta_{max} = 2\text{cm}$. The management action therefore represents an intervention to correct for an excessive displacement at a location, thereby avoiding high consequences resulting from the displacement, while paying a small cost for the corrective action.

Three types of measurements are available to support decision-making. Measurements of type a evaluate the horizontal displacement in one location, with negligible measurement noise. Type b measures evaluate the vertical displacement. However, these vertical measures have a significant spatially correlated bias. This bias is modeled as a Gaussian process with zero mean and covariance:

$$k_{bias}(x, x') = \sigma_{bias}^2 \exp\left[-\frac{(x-x')^2}{2{\lambda_{bias}}^2}\right] \tag{15}$$

where $\sigma_{bias} = 100\text{cm}$ and $\lambda_{bias} = 20\text{m}$. This covariance function is used to define $\boldsymbol{\mu}_\epsilon$ and $\boldsymbol{\Sigma}_\epsilon$ when type b measures are included. Finally, type c measures are calibration measures, which determine the bias of type b measures directly. Measurements of types a and b have a cost of 1, while type c calibrations have a cost of 4. Type a, b, and c measures may be selected along the domain with a minimum spacing of 0.17m.

To support system management, measurements are optimized via the forward greedy approach using the VoI metric (in this example there are too many potential sensor placements considered for the reverse approach to be implemented in a reasonable amount of time). Furthermore, the surrogate conditional entropy metric discussed in Section 4.2 is used, together with a forward greedy optimization approach based on combining this metric with the VoI, as in Algorithm 3, to provide an alternative optimal measurement scheme which seeks to avoid pitfalls related to the non-submodularity of the VoI. For comparative purposes, the conditional entropy metric alone (optimized using a forward greedy algorithm) as well as a genetic algorithm optimizing VoI are also used to propose optimal measurement schemes. Note that for the conditional entropy metric, the entropy of all random variables, including the bias of the type b measures, is considered. For the genetic algorithm, as in the example of Section 5.2, two separate runs are performed: one which is constrained to match the runtime of the heuristic approach and find solutions of various sizes, and one which is unconstrained, allowing it to determine the size and composition of the optimal measurement set without being subject to time constraints. In this case, with the variables to be optimized representing whether to include each possible measurement location and type in the selected set, the maximum number of generations is 3.6×10^4. All other implementation details are the same as in Section 5.2.

Fig. 4a shows the measurements of various types chosen as the optimal set via the approaches considered. In this problem, the forward greedy approach using VoI falls into a pattern early on of selecting type a measurements, as these measures provide some value in and of themselves, while type b measures are only valuable when combined with type c calibrations. Thus, if left to run until completion, this approach would select all type a measures until these are exhausted before moving on to select type b and c measures. However, at this point the cost incurred to obtain all type a measures would far outweigh any value the type b or c measures might now provide. The forward greedy approach based on VoI therefore performs relatively poorly in this problem. Similarly, the forward greedy approach based on the conditional entropy metric alone only focuses on type a and c measures, to the neglect of type b measures. Because this metric does not involve any information related to the decision-making context for the problem, it is unaware of, and therefore cannot prioritize, the relationship between type b and c measures in determining the vertical displacement and therefore in guiding management activities. For this reason, despite not being confined to type a measures only, the entropy-based approach underperforms approaches which make use of the VoI. This is consistent with previous results that show that entropy alone is often an insufficient criterion for selection optimal sensor placement to support decision-making [7].

The proposed heuristic approach combining the conditional entropy surrogate metric with VoI follows the forward greedy approach for VoI until about the tenth sensor is placed. At this point, the additional benefits of type a measures fall below their costs, and the conditional entropy metric is used, leading to the selection of type c measures. Switching back to the VoI metric, the benefits of combining type b and c

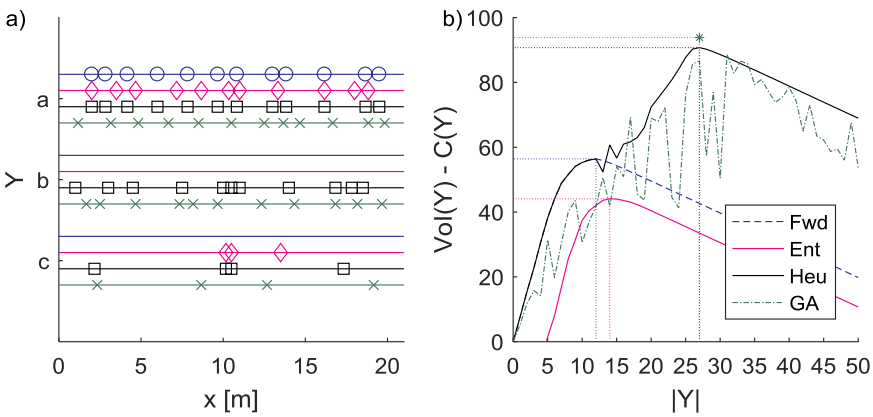

Fig. 4. a) Selected optimal measurement sets by forward greedy optimization of VoI (Fwd, ○), greedy optimization using a submodular surrogate (Heu, □), forward greedy optimization of conditional entropy (Ent, ◊), and generic algorithm optimization of VoI (GA, ×), with position indicated on the horizontal axis and measurement type indicated by three groupings along the vertical axis (results for each algorithm are slightly vertically offset to improve clarity); b) Net VoI versus number of sensors for each approach is indicated by solid lines, and the peaks of these curves, corresponding to the sets plotted in part (a), are indicated with dotted lines. Note that for conditional entropy optimization, the VoI of measurement sets selected by the greedy optimization has been evaluated and plotted. Also note that for the genetic algorithm, optimized sets of various sizes resulting from constrained runs are represented by lines for comparison purposes, while the overall solution obtained for a separate unconstrained run (the results of which are depicted to the left) is indicated by the green star. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

measures are identified, which greatly increases the net VoI. Therefore, in this problem, the heuristic approach using the conditional entropy metric serves as an improvement over the standard forward greedy optimization using the VoI only, and can be implemented in a reasonable amount of time, unlike the reverse greedy approach.

Finally, compared to the stochastic approach of the genetic algorithm solver, the proposed heuristic represents a deterministic approximation. Thus, when constrained to use the same runtime, the genetic algorithm sometimes gives a solution near that of the heuristic method, and occasionally even outperforms the heuristic, but also sometimes performs relatively poorly (see green line in Fig. 4b). When allowed to perform optimization without a time constraint, the genetic algorithm does ultimately arrive at a more optimal solution than the proposed heuristic, outperforming the latter by 3% in terms of net VoI (see green star in Fig. 4b). However, this comes at the cost of a runtime which is greater in this problem by a factor of four. The ability of this efficient heuristic based on forward greedy optimization to achieve similar results compared to the genetic algorithm can have beneficial implications for other, larger problems in combinatorial optimization for sensor placement and scheduling. Overall, while the genetic algorithm's stochastic approach is less constrained in how it searches the solution space, the proposed heuristic is more efficient while achieving a solution of similar quality (at least for the examples presented above). However, it must also be noted that, like the genetic algorithm approach, no general guarantees of optimality can be provided for this heuristic approach, and there may be certain circumstances in which its performance will be rather poor.

6. Conclusions

In this paper, the problem of optimal sensor placement and scheduling is examined. While the VoI metric directly quantifies the benefits of collected information in supporting management decision-making, greedy optimization based on this metric can be hampered by its lack of submodularity, as has been demonstrated through several examples.

A practical approach to overcoming these difficulties is to apply multiple optimization approaches to identify different potential optimal measurement sets and then compare the VoI of these sets to determine the best among them. We have presented the forward and reverse greedy approximate optimization approaches, as applied directly to the VoI metric, and a submodular surrogate metric based on conditional entropy. We have investigated their performance, and compared it with that of state-of-the-art optimization methods such as genetic algorithms. While none of these approaches is guaranteed to deliver the true optimal solution, by comparing multiple potential solutions in this manner there is a greater chance to avoid choosing a highly suboptimal solution arrived at using only one method. Furthermore, it should be noted that VoI is upper-bounded by the value of perfect or complete information (VoPI) for a problem, assessed as the VoI of obtaining all possible measures in the system, i.e. as $\mathrm{VoI}(\mathcal{Y})$. Assessing the VoPI at the beginning of the optimization allows for the VoI resulting from the chosen optimization method to be compared to the VoPI, to estimate the potential maximum improvement related to extensive monitoring.

Greedy optimization approaches, although efficient, do not assess the value of multiple measurements taken together. However, this shortcoming results directly from the efficiency of the algorithm and cannot be overcome without sacrificing this efficiency to some degree. For example, generalizations of a greedy approach which consider sets of q additional measurements at each iteration could detect sets of this size which provide a high VoI when taken together. However, the computational demand of such an approach grows rapidly with q, and there is always the possibility that it is a set of $q + 1$ measurements taken together which provides the significant value. On the other hand, it may be possible to employ further heuristics to assist this search process more efficiently. An efficient greedy approach is an important tool for optimizing sensing via the VoI but must be informed by knowledge of its limitations and how to overcome these when necessary.

Acknowledgements

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”.

Appendix A: Gaussian Process Spatio-Temporal Models

The examples of this paper make use of Gaussian process spatio-temporal random field models to describe the variables which define system performance. Gaussian processes represent a generalization of the multivariate Gaussian distribution to a continuous domain [24]. A Gaussian

process spatio-temporal model is denoted as:

$$\mathrm{f}(\mathbf{x}, t) \sim \mathcal{GP}[\mu(\mathbf{x}, t), \mathrm{k}(\mathbf{x}, t, \mathbf{x}', t')] \tag{16}$$

where $\mu(\mathbf{x},t)$ is the mean function, defining the mean of the random field at spatial coordinate $\mathbf{x}$ and time t, and $\mathrm{k}(\mathbf{x},t,\mathbf{x}',t')$ is the covariance function, defining the covariance between random variables associated with locations $\mathbf{x}$ and $\mathbf{x}'$ and times t and t'. For a given set of spatial coordinates and a management time duration, this model defines a joint Gaussian distribution for the vector of random field variables affecting the system:

$$\mathbf{f} \sim \mathcal{N}(\boldsymbol{\mu}_F, \Sigma_F) \tag{17}$$

where $\boldsymbol{\mu}_F$ is the mean vector and $\boldsymbol{\Sigma}_F$ is the covariance matrix for the joint distribution.

Observations of these random field variables under a given measurement scheme Y are described as noisy measurements of linear combinations of these variables as follows:

$$y = \mathbf{R}_Y\mathbf{f} + \epsilon \quad \epsilon \sim \mathcal{N}(\boldsymbol{\mu}_\epsilon, \boldsymbol{\Sigma}_\epsilon) \tag{18}$$

where matrix $\mathbf{R}_Y$ encodes the relationships between individual measures and which random variables (or linear combinations of variables) are to be observed. Note that $\mathbf{R}_Y$ must be formulated appropriately so as to prohibit the measurement of future random variables. The measurement noise is described by ϵ, which is assumed to have a joint Gaussian distribution with mean $\boldsymbol{\mu}_\epsilon$ and covariance $\boldsymbol{\Sigma}_\epsilon$. This covariance matrix can include correlated errors between measurements, representing systematic bias and/or sensor drift. The above definition gives the measurement vector $\mathbf{y}$ a Gaussian distribution with mean $\boldsymbol{\mu}_Y = \mathbf{R}_Y\boldsymbol{\mu}_F + \boldsymbol{\mu}_\epsilon$ and covariance $\boldsymbol{\Sigma}_Y = \mathbf{R}_Y\boldsymbol{\Sigma}_F\mathbf{R}_Y^{\mathrm{T}} + \boldsymbol{\Sigma}_\epsilon$.

A large degree of mutual correlation between Gaussian errors is equivalent to uncorrelated errors combined with an unknown systematic bias common to all errors. To illustrate this, consider the following example where the error ϵ for s measures is the sum of an unknown systematic bias, with standard deviation σ_b, across all measurements and independent noise errors, with unitary standard deviation, for each measurement:

$$\epsilon = \epsilon_{\mathrm{bias}}\cdot\mathbf{1}_{s,1} + \epsilon_{\mathrm{noise}} \quad \epsilon_{\mathrm{bias}} \sim \mathcal{N}(\mathbf{0}, \sigma_b^2) \quad \epsilon_{\mathrm{noise}} \sim \mathcal{N}(\mathbf{0}_{s,1}, \mathbf{I}) \tag{19}$$

Where $\mathbf{1}_{a,c}$ and $\mathbf{0}_{a,c}$ indicates the matrix (or vector) of ones and zeros respectively, of dimension $a \times c$. Using the standard rules for adding independent Gaussian random variables, the mean and covariance of the overall error can be computed as follows:

$$\boldsymbol{\mu}_\epsilon = \mathbf{0}_{s,1} \quad \boldsymbol{\Sigma}_\epsilon = \sigma_b^2\mathbf{1}_{s,s} + \mathbf{I} \tag{20}$$

Corresponding to a correlation coefficient $\rho = 1/(1 + \sigma_b^{-2})$. For example, ρ is 91% if σ_b is 10, that is if the uncertainty in the bias is 10 times higher than the uncorrelated noise. Thus, the combination of a random independent noise and a systematic bias can be expressed equivalently using a multivariate Gaussian distribution.

Given a set of measurements $\mathbf{y}$ of the random field, the prior underlying random field model for the system is updated as follows:

$$\mathbf{f}|\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}_{F|\mathbf{y}}, \boldsymbol{\Sigma}_{F|Y}) \tag{21}$$

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 $\boldsymbol{\Sigma}_{F|Y} = \boldsymbol{\Sigma}_F - \mathbf{R}_Y\boldsymbol{\Sigma}_F\boldsymbol{\Sigma}_Y^{-1}\boldsymbol{\Sigma}_F\mathbf{R}_Y^{\mathrm{T}}$. Note that while $\boldsymbol{\mu}_{F|\mathbf{y}}$ is a function of the vector $\mathbf{y}$ of measurements, $\boldsymbol{\Sigma}_{F|Y}$ is a function of the measurement scheme Y only.

References

[1] Raiffa H, Schlaifer R. Applied statistical decision theory. Cambridge, Massachusetts, USA: Harvard University Press; 1961.
[2] Pozzi M, Der Kiureghian A. Computation of lifetime value of information for monitoring systems. in *Proceedings of the 8th International Workshop on Structural Health Monitoring*. Stanford University; 2011.
[3] Straub D. Value of information analysis with structural reliability methods. Struct Saf 2014;49:75–86.
[4] Memarzadeh M, Pozzi M. Integrated inspection scheduling and maintenance planning for infrastructure systems: integrated inspection scheduling and maintenance planning. Comput-Aid Civil Infrastruct Eng Oct. 2015. p. n/a-n/a.
[5] 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.
[6] Bellman R. Dynamic programming. Princeton, NJ: Princeton Univ. Pr; 1984.
[7] Malings C, Pozzi M. Conditional entropy and value of information metrics for optimal sensing in infrastructure systems. Struct Saf 2016;60:77–90.
[8] Malings C, Pozzi M. Value-of-information in spatio-temporal systems: sensor placement and scheduling. Reliab Eng Syst Saf 2018;172:45–57.
[9] Krause A. Optimizing Sensing: Theory and Applications 15213. Pittsburgh, PA: Carnegie Mellon University; 2008.
[10] Krause A, McMahon HB, Guestrin C, Gupta A. Robust submodular observation selection. J Mach Learn Res 2008;9:2761–801.
[11] Krause A, Singh A, Guestrin C. Near-optimal sensor placements in Gaussian processes: theory, efficient algorithms and empirical studies. J Mach Learn Res 2008;9:235–84.
[12] Malings C, Pozzi M. Value of information for spatially distributed systems: application to sensor placement. Reliab Eng Syst Saf Oct. 2016;154:219–33.
[13] Nemhauser G, Wolsey L, Fisher M. An analysis of the approximations for maximizing submodular set functions. Math Program 1978;14:265–94.
[14] Clark A, Alomair B, Bushnell L, Poovendran R. Submodularity in dynamics and control of networked systems. Cham: Springer International Publishing; 2016.
[15] Cressie NAC, Wikle CK. Statistics for spatio-temporal data. Hoboken, N.J: Wiley; 2011.
[16] Papadimitriou C. Optimal sensor placement methodology for parametric identification of structural systems. J Sound Vibr Dec. 2004;278(4–5):923–47.
[17] Krause A, Guestrin C. Optimal nonmyopic value of information in graphical models –efficient algorithms and theoretical limits. in *Proceedings of the 19th International Joint Conference on Artificial Intelligence*. 2005.
[18] Krause A, Guestrin C. Optimal value of information in graphical models. J Artif Intell Res 2009;35:557–91.
[19] Malings C, Pozzi M. Sensor network optimization using bayesian networks, decision graphs, and value of information. in *Proceedings of the 12th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP12)*. 2015.
[20] Fujishige S. Polymatroidal dependence structure of a set of random variables. Inf Control 1978;39:55–72.
[21] Cover T, Thomas J. Elements of information theory. 2nd ed Hoboken, New Jersey, USA: John Wiley & Sons, Inc.; 2006.
[22] 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. 2016.
[23] Guo HY, Zhang L, Zhang LL, Zhou JX. Optimal placement of sensors for structural health monitoring using improved genetic algorithms. Smart Mater Struct Jun. 2004;13(3):528–34.
[24] Rasmussen CE, Williams CKI. Gaussian processes for machine learning. Cambridge, Mass: MIT Press; 2006.