2015 IEEE 54th Annual Conference on Decision and Control (CDC)
December 15-18, 2015. Osaka, Japan

Resilient Observation Selection in Adversarial Settings

Aron Laszka, Yevgeniy Vorobeychik, and Xenofon Koutsoukos
Vanderbilt University

Abstract— **Monitoring large areas using sensors is fundamental in a number of applications, including electric power grid, traffic networks, and sensor-based pollution control systems. However, the number of sensors that can be deployed is often limited by financial or technological constraints. This problem is further complicated by the presence of strategic adversaries, who may disable some of the deployed sensors in order to impair the operator's ability to make predictions. Assuming that the operator employs a Gaussian-process-based regression model, we formulate the problem of attack-resilient sensor placement as the problem of selecting a subset from a set of possible observations, with the goal of minimizing the uncertainty of predictions. We show that both finding an optimal resilient subset and finding an optimal attack against a given subset are NP-hard problems. Since both the design and the attack problems are computationally complex, we propose efficient heuristic algorithms for solving them and present theoretical approximability results. Finally, we show that the proposed algorithms perform exceptionally well in practice using numerical results based on real-world datasets.**

I. Introduction

Our ability to dynamically control any system hinges on having accurate information about its evolving state, obtained through persistent system monitoring. In many applications, such as electric power grids or traffic networks, the system to be monitored can extend over a vast area, with many possible points of observation. Although these areas can be very large, the number of sensors that can be deployed is limited by financial and/or technological constraints. Consequently, we are faced with a problem of finding locations for placing a limited number of sensors so as to minimize our posterior uncertainty about the quantities being monitored. Due to its importance, this problem of sensor placement (or, more generally, observation/feature selection) and associated predictions about unobserved state variables has received considerable attention, particularly when variables of interest are modeled using a Gaussian process regression. For example, Gaussian process regression models have been successfully applied to a wide range of problem, such as traffic volume forecasting [20], [5], spatial modeling of extreme snow depth [2], wind power forecasting [12], estimation of water chlorophyll concentration [1], and spectrum sensing [17].

Despite considerable attention, however, little past research has investigated the problem of sensor placement or observation selection in adversarial settings. For example, while this problem was briefly considered by Krause et al. [14], their approach requires one to define different adversarial objectives for every possible attack, clearly not a scalable proposition as we increase the attacker's capabilities (that is, increase the number of sensors that can be attacked). In many applications, however, the sensors are not physically well protected, since they can be widely geographically distributed. Moreover, cyber attacks on sensors may circumvent physical security, particularly since in many applications sensor readings are transmitted over the internet [19], [10]. Hence, an adversary may mount a denial-of-service (DoS) attack, such as wireless jamming, or flooding the VPN's external interface to disrupt SCADA communications in the electric power grid [10], potentially causing significant financial losses or even physical damage.

We address the problem of finding a sensor deployment that is resilient against denial-of-service attacks on a subset of sensors. More formally, we consider the problem of finding an observation selection scheme which minimizes uncertainty, accounting for a DoS attack on a subset of sensors. We show that, in general, both the problem of optimal DoS attack and that of resilient observation selection are NP-hard. On the positive side, we develop efficient approximation algorithms for both problems, and exhibit approximability guarantees for the resilient observation selection algorithm. Additionally, we develop efficient optimal attack and resilient observation selection algorithms in the special case when the covariance matrix among variables has a tree structure. Clearly, by selecting a subset of observations that differs from that optimal in the "no-attack" case, we may increase resilience, but at the cost of increased uncertainty about unobserved variables under normal system operations. We use real data to demonstrate that the increase in uncertainty under normal operations due to resilient observation selection scheme is quite small, whereas the resulting system is substantially more resilient to attacks than the baseline.

The remainder of this paper is organized as follows. In Section II, we give a brief summary of related work. In Section III, we formulate the resilient observation selection problem. In Section IV, we present theoretical results. In Section V, we present numerical results. Finally, in Section VI, we conclude the paper.

II. Related Work

A. Observation Selection

The problem of selecting a subset of observation variables has been studied in the context of both linear regression and Gaussian process regression. Here, we summarize the most relevant previous work.

Das and Kempe study the problem of selecting a subset of random variables to observe from a large set, in order to obtain the best linear predictor of another variable of interest. In [6], they introduce exact and approximation

978-1-4799-7886-1/15/$31.00 ©2015 IEEE

algorithms for several special cases of practical interest. In [7], they obtain strong approximation guarantees on widely-used greedy heuristics by introducing the submodularity ratio, which captures how close to submodular an objective function is.

Krause et al. study the problem of selecting observations which perform well when evaluated with an objective function chosen by an adversary [13], [14]. They present an algorithm, called the Submodular Saturation algorithm, which has strong theoretical approximation guarantees when the objective functions exhibit submodularity.

Krause et al. also study the problem in non-adversarial settings in [15], where they consider the problem of maximizing the mutual information between the chosen variables and the variables that are not selected. The authors show that the problem is NP-hard, propose efficient greedy heuristics, and finally consider random node failures and uncertainties in the model.

B. Sensor Placement

The problem of sensor placement has been extensively studied, especially in the context of sensor networks (for a comprehensive review of the literature, we refer the reader to the survey of Younis and Akkaya [21]). However, most of the previous work on sensor placement focuses primarily on area coverage and network connectivity. For example, Dhillon and Chakrabarty [9] propose two algorithms for the efficient placement of sensors. These algorithms maximize either the average or the minimum coverage of grid points, respectively, under the constraints of uncertain sensor readings and terrain properties.

Nevertheless, a number of papers have considered resilient node placement; however, most of these are either concerned with tolerance against random faults or with resilient network structures. For example, Ishizuka and Aida [11] study the tolerance of various non-deterministic node placement strategies against non-adversarial battery exhaustion and random failures; however, the authors do not consider deliberate attacks, which cause worst-case failures. As another example, Bredin et al. [3] consider the problem of deploying a sensor network to guarantee a specified level of k-connectivity. Such approaches, which consider the connectivity of the network, are complementary to our work, which focuses on the information gathered by the sensors.

Finally, several papers are concerned with sensor placement for specific applications. For example, Carr et al. [4] study the problem of placing sensors in municipal water networks to detect contaminants. The authors formulate a series of sensor placement problems as mixed-integer programs, whose objective coefficients are not known with certainty.

III. Problem Formulation

Notation

Let σ_Y^2 denote the variance of a random variable Y, $\sigma_{Y|\mathcal{S}}^2$ denote the variance of a random variable Y given the values of variables in a set $\mathcal{S}$, Σ_{YX} denote the covariance between random variables Y and X, $\mathbf{\Sigma}_{\mathcal{ST}}$ denote the submatrix formed by the rows in set $\mathcal{S}$ and the columns in set $\mathcal{T}$.

A. Gaussian Process Regression

Here, we summarize the basics of Gaussian process regression that are necessary for our paper. For a comprehensive discussion of Gaussian processes in machine learning, we refer the reader to [18].

We model both the possible observations and the predictor as random variables, whose joint distribution is a Gaussian process. As an example, in the traffic monitoring application, each random variable could represent the traffic density on a given road segment. Then, given observed values $\boldsymbol{x}_{\mathcal{S}}$ at set $\mathcal{S}$, the predictor variable Y follows a Gaussian distribution $\mathcal{N}\left(\mu_{Y|\mathcal{S}}, \sigma_{Y|\mathcal{S}}^2\right)$, where

$$\mu_{Y|\mathcal{S}} = \mu_Y + \mathbf{\Sigma}_{Y\mathcal{S}}\mathbf{\Sigma}_{\mathcal{SS}}^{-1}\left(\boldsymbol{x}_{\mathcal{S}} - \boldsymbol{\mu}_{\mathcal{S}}\right) \quad (1)$$

$$\sigma_{Y|\mathcal{S}}^2 = \sigma_Y^2 - \mathbf{\Sigma}_{Y\mathcal{S}}\mathbf{\Sigma}_{\mathcal{SS}}^{-1}\mathbf{\Sigma}_{\mathcal{S}Y} \ , \quad (2)$$

where $\mathbf{\Sigma}$ is the (prior) covariance matrix of all the variables, while μ_Y and $\boldsymbol{\mu}_{\mathcal{S}}$ are the (prior) means of the variables.

The value of $\sigma_{Y|\mathcal{S}}^2$, which we will call the *posterior variance of* Y, is of interest to us for multiple reasons. First, suppose that we have to predict the value of Y based on observations at $\mathcal{S}$ such that we minimize the mean squared error of the prediction. Then, it is easy to see that the error is minimized at $\mu_{Y|\mathcal{S}}$ and the minimal error value is $\sigma_{Y|\mathcal{S}}^2$. Second, suppose that we are interested in the uncertainty of Y given observations at $\mathcal{S}$, which we measure using entropy. Then, it is well-known that this entropy is logarithmically proportional to $\sigma_{Y|\mathcal{S}}^2$; more specifically, its value is $\frac{1}{2}\ln\left(2\pi e\sigma_{Y|\mathcal{S}}^2\right)$.

Finally, observe that the value of posterior variance $\sigma_{Y|\mathcal{S}}^2$ depends only on the set of observations $\mathcal{S}$ and the (prior) covariances $\mathbf{\Sigma}$, but not on the actual observed values $\boldsymbol{x}_{\mathcal{S}}$. This observation is important, since it will allow us to define the objective of the observation selection problem as a set function of $\mathcal{S}$ parametrized by $\mathbf{\Sigma}$.

B. Observation Selection

We now introduce the observation selection problem, which was studied in [6] and [7].[1] Note that we will sometimes refer to this problem as the *non-resilient* observation selection problem to emphasize the distinction between this problem and the resilient observation selection problem, which we will define in the following section.

In the observation selection problem, our goal is to select an N-sized set of observations $\mathcal{S}$ from a set of possible observations $\mathcal{V}$ so that the posterior variance of a target variable Y that we wish to predict is minimized. Formally, given N, $\mathcal{V}$, and Y, we have to find

$$\operatorname*{argmin}_{\mathcal{S}\subseteq\mathcal{V}:|\mathcal{S}|=N} \sigma_{Y|\mathcal{S}}^2 \ . \quad (3)$$

[1]Note that, in [6] and [7], the problem was motivated by a linear regression based on features that were modeled as random variables.

To study the computational complexity of the observation selection problem, we can reformulate it as a decision problem as follows.

Definition 1: Observation Selection Problem [decision version]: Given a predictor variable Y, a set of variables $\mathcal{V}$, the covariance matrix $\mathbf{\Sigma}$ of these variables, a selection size $N \in \mathbb{N}_+$, and a threshold variance $T \in \mathbb{R}_+$, determine if

$$T \geq \min_{\mathcal{S} \subseteq \mathcal{V}:|\mathcal{S}|=N} \sigma^2_{Y|\mathcal{S}} \ . \tag{4}$$

Unfortunately, the observation selection problem is NP-hard in general [6]. This is true even for such restricted cases as deciding whether zero posterior variance is attainable [8].

C. Resilient Observation Selection

In this paper, we study the problem of resilient observation selection in adversarial environments. More specifically, our goal is to find efficient algorithms for selecting a set of observations that minimize the posterior variance of the predictor variable, given that some of the observations will be removed by a strategic adversary.

We assume that a set of possible observations (formally, a set of random variables) are given to a designer, who can select N of these. In practice, the set of possible observations can model, for example, the set of possible locations at which sensors can be placed, and the cardinality constraint is due to budget or technological constraints. We assume that a strategic adversary can remove K of the selected observations.[2] In practice, removing a node can model all forms of denial-of-service type attacks, such as physical destruction, wireless jamming, or battery exhaustion, and the cardinality constraint is again due to budget or technological constraints, now on the adversary. The goals of the designer and the adversary are opposed: the designer's goal is to minimize the posterior variance of the predictor variable given the set of selected but non-removed observations, while the adversary's goal is to maximize the same value.

Formally, the resilient observation selection problem is defined as

$$\operatorname*{argmin}_{\mathcal{S} \subseteq \mathcal{V}:|\mathcal{S}|=N} \left(\max_{\mathcal{A} \subseteq \mathcal{S}:|\mathcal{A}|=K} \sigma^2_{Y|(\mathcal{S} \setminus \mathcal{A})} \right) \ , \tag{5}$$

where $\mathcal{V}$ is the set of possible observations, N is the number of observations selected, and K is the number of observations removed by the adversary. Notice that the order of $argmin$ and max in the above formulation models the order in which the designer and the adversary make their choices. Specifically, we assume that the defender moves first, placing the sensors $\mathcal{S}$, while the adversary subsequently observes the placement of sensors and chooses to remove a subset $\mathcal{A}$ of them (e.g., through a denial-of-service attack).

IV. Analysis

In this section, we present computational complexity results, propose efficient heuristic and approximation algorithms, and study a special case of the resilient selection problem. First, in Section IV-A, we consider the subproblem of finding an optimal attack against a given selection. Then, in Section IV-B, we study the problem of resilient observation selection. Finally, in Section IV-C, we focus on a special case of the problem, for which the optimal solution can be found in polynomial time.

A. Attack Problem for a Given Selection

We begin our analysis with studying the problem faced by the adversary: finding an optimal attack against a given selection of observations. This problem is of interest to us because finding an optimal attack is necessary to quantifying the resilience of a given observation subset. Formally, for a given selection $\mathcal{S}$, the attacker aims to solve

$$\max_{\mathcal{A} \subseteq \mathcal{S}:|\mathcal{A}|=K} \sigma^2_{Y|(\mathcal{S} \setminus \mathcal{A})} \ . \tag{6}$$

To study the computational complexity of this problem, we define a decision version of it as follows.

Definition 2: Optimal Attack Problem [decision version]: Given a predictor variable Y, a set of variables $\mathcal{S}$, the covariance matrix $\mathbf{\Sigma}$ of these variables, an attack size $K \in \mathbb{N}$, and a threshold variance $T \in \mathbb{R}_+$, determine if

$$T \leq \max_{\mathcal{A} \subseteq \mathcal{S}:|\mathcal{A}|=K} \sigma^2_{Y|(\mathcal{S} \setminus \mathcal{A})} \ . \tag{7}$$

The following theorem shows that, in general, finding an optimal attack is a computationally challenging problem.

Theorem 1: The Optimal Attack Problem is NP-hard.

The proof of Theorem 1 can be found in [16].

Some of the most commonly used heuristics for the (non-resilient) selection problem are greedy heuristics. In the following definition, we formulate a straightforward greedy algorithm for the attack problem.

Definition 3: Greedy Algorithm for Finding an Attack:

1: $\mathcal{A} \leftarrow \emptyset$
2: **while** $|\mathcal{A}| < K$ **do**
3: $\quad X^* \in \operatorname{argmax}_{X \in \mathcal{S} \setminus \mathcal{A}} \sigma^2_{Y|(\mathcal{S} \setminus (\mathcal{A} \cup \{X\}))}$
4: $\quad \mathcal{A} \leftarrow \mathcal{A} \cup \{X^*\}$
5: **end while**
6: **return** $\mathcal{A}$

Unfortunately, the output of the greedy algorithm can be arbitrarily worse than the optimal solution.

Proposition 1: For any $\delta > 0$, there exist an instance of the optimal attack problem such that $\sigma^2_{Y|\textsc{Greedy}} \leq \delta \cdot \sigma^2_{Y|\textsc{Optimal}}$, where GREEDY is the output of the greedy algorithm and OPTIMAL is an optimal solution.

The proof of Proposition 1 can be found in [16].

Nonetheless, we will see in Section V-B that the greedy algorithm performs extremely well in practice.

B. Resilient Observation Selection Problem

1) Computational Complexity: Now, we tackle our main problem of resilient observation selection. To study the computational complexity of this problem, we formulate it as a decision problem as follows.

Definition 4: Resilient Observation Selection Problem [decision version]: Given a predictor variable Y, a set of

[2]Note that we can assume $K < N$. Otherwise, the problem is trivial.

variables $\mathcal{V}$, the covariance matrix $\mathbf{\Sigma}$ of these variables, a selection size $N \in \mathbb{N}$, an attack size $K \in \mathbb{N}$, and a threshold variance $T \in \mathbb{R}_+$, determine if

$$T \geq \min_{\mathcal{S}\subseteq\mathcal{V}:|\mathcal{S}|=N} \left(\max_{\mathcal{A}\subseteq\mathcal{S}:|\mathcal{A}|=K} \sigma^2_{Y|(\mathcal{S}\setminus\mathcal{A})} \right) . \quad (8)$$

Notice that the non-resilient selection problem, which we know to be NP-hard, is the special case of $K = 0$. The following proposition shows that the resilient selection problem is also NP-hard for any fixed $K > 0$.

Proposition 2: The Resilient Observation Selection Problem is NP-hard for any fixed attack size K.

The proof of Proposition 2 can be found in [16].

2) Greedy Algorithm for Resilient Selection: In practice, the non-resilient observation selection problem is often solved using greedy algorithms, which are indeed very good heuristics [7]. These algorithms start with an empty set and add observations to this set one-by-one, always picking the one that decreases the objective function (i.e., the posterior variance of Y) the most.

Unfortunately, we cannot directly apply this approach to the resilient selection problem, since its objective function is not well-defined (or zero) as long as fewer than $K+1$ nodes are selected. To address this problem, we select the first $K+1$ observations in a single step, picking the $(K+1)$-subset that maximally decreases the objective function. The latter subproblem, however, poses its own challenge: if we were to exhaustively search all $(K+1)$ subsets of observations, we would face running time which is exponential in K. Fortunately, we can use the following simple observation to find the optimal subset efficiently: if the designer selects only $K+1$ observations, then there will only be a single observation that survives the attack. Furthermore, the adversary will always leave intact the observation which carries the least information regarding the predictor. In formal terms, if $|\mathcal{S}| = K+1$, then

$$\max_{\mathcal{A}\subseteq\mathcal{S}:|\mathcal{A}|=K} \sigma^2_{Y|(\mathcal{S}\setminus\mathcal{A})} = \max_{X\in\mathcal{S}} \sigma^2_{Y|X} . \quad (9)$$

Consequently, we can find an optimal $(K+1)$-set by selecting the $K+1$ observations that each carry the most information regarding the predictor.

Based on the above idea, we can define a greedy algorithm for resilient observation selection as follows.

Definition 5: Greedy Algorithm for Resilient Observation Selection:

$\mathcal{S} \leftarrow \emptyset$
while $|\mathcal{S}| < K+1$ **do**
  $X^* \in \operatorname{argmin}_{X\in(\mathcal{V}\setminus\mathcal{S})} \sigma^2_{Y|X}$
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X^*\}$
end while
while $|\mathcal{S}| < N$ **do**
  $X^* \in \operatorname{argmin}_{X\in(\mathcal{V}\setminus\mathcal{S})} \max_{\mathcal{A}\in(\mathcal{S}\cup\{X\}):|\mathcal{A}|=K} \sigma^2_{Y|((\mathcal{S}\cup\{X\})\setminus\mathcal{A})}$
  $\mathcal{S} \leftarrow \mathcal{S} \cup \{X^*\}$
end while
return $\mathcal{S}$

Note that, in the second loop, the above algorithm uses exhaustive search to find the optimal attack. In practice, we can use the greedy algorithm introduced in Section IV-A instead, since its performance is almost optimal in practice (see Section V-B).

Our next step is to establish an approximation bound on the performance of the above greedy resilient observation selection algorithm. For this purpose, we first reformulate our selection problem as a maximization problem by considering the decrease in posterior variance to be our objective function. Formally, the resilient selection problem can be formulated as

$$\operatorname*{argmax}_{\mathcal{S}\subseteq\mathcal{V}:|\mathcal{S}|=N} \left(\min_{\mathcal{A}\subseteq\mathcal{S}:|\mathcal{A}|=K} \sigma^2_Y - \sigma^2_{Y|(\mathcal{S}\setminus\mathcal{A})} \right) . \quad (10)$$

Note that this problem is equivalent to Equation (5).

We begin by generalizing the concept of submodularity ratio, which was introduced by Das and Kempe in [7].[3] This generalization is necessary, since applying the original, non-resilient definition to our adversarial objective function (i.e., $\min_{\mathcal{A}\subseteq\mathcal{S}:|\mathcal{A}|=K} \sigma^2_Y - \sigma^2_{Y|(\mathcal{S}\setminus\mathcal{A})}$) would result in no guarantees.[4]

Definition 6: Resilient Submodularity Ratio: Given a non-negative non-decreasing set function f, the *resilient submodularity ratio* of f with respect to a set $\mathcal{U}$ and parameters N and K is

$$\gamma_{\mathcal{U},N,K}(f) = \min_{\substack{\mathcal{L},\mathcal{S}:\\ \mathcal{L}\subseteq\mathcal{U},|\mathcal{L}|\geq K,\\ |\mathcal{S}|\leq N,\mathcal{S}\cap\mathcal{L}=\emptyset}} \frac{\sum_{X\in\mathcal{S}} [f(\mathcal{L}\cup\{X\}) - f(\mathcal{L})]}{f(\mathcal{L}\cup\mathcal{S}) - f(\mathcal{L})} . \quad (11)$$

Note that, by comparing the above definition with the definition of the (non-resilient) submodularity ratio in [7], we can see that the non-resilient ratio is a special case of the resilient ratio with $K = 0$.

Using the notion of resilient submodularity ratio, we can now prove the following approximation bound on our greedy resilient observation selection algorithm.

Theorem 2: Let $f(\mathcal{S}) = \min_{\mathcal{A}\subseteq\mathcal{S}:|\mathcal{A}|=K} \sigma^2_Y - \sigma^2_{Y|(\mathcal{S}\setminus\mathcal{A})}$, let $\mathcal{S}^G$ be the set selected by the greedy algorithm, and let $\mathrm{OPT} = \max_{\mathcal{S}\subseteq\mathcal{V}:|\mathcal{S}|=N} f(\mathcal{S})$. Then,

$$f(\mathcal{S}^G) \geq \mathrm{OPT} - \mathrm{OPT} \cdot \left(1 - \frac{\gamma_{\mathcal{S}^G,N,K}(f)}{N}\right)^{N-K} . \quad (12)$$

The proof of Theorem 2 can be found in [16].

Note that, if f is a submodular function, then the submodularity ratio $\gamma_{\mathcal{U},N,K}(f)$ is equal to 1 by definition (regardless of the parameters), which results in the lower bound

$$\mathrm{OPT} - \mathrm{OPT} \cdot \left(1 - \frac{1}{N}\right)^{N-K} . \quad (13)$$

In practice, the objective function is usually close to submodular, which means that the output of the greedy algorithm can be expected to be very close to optimal. This is confirmed by our numerical results in Section V.

[3]Recall from Section II that the submodularity ratio measures how close a function is to being submodular, and it was used in [7] to establish approximation guarantees on the (non-resilient) greedy algorithm.

[4]The submodularity ratio of our objective function would always be zero.

C. Special Case: Tree Covariance Graphs

In Proposition 2, we used the NP-hardness of the non-resilient selection problem to show that the resilient selection problem is also NP-hard in general. However, it was shown in [6] that the non-resilient selection problem can be solved in polynomial time for certain special cases. Consequently, the following question arises naturally: can the resilient observation selection problem also be solved efficiently for these special cases?

In this subsection, we show that this is indeed possible for the special case of *tree covariance graphs*. First, to facilitate the characterization of this special case, we introduce the notion of covariance graphs.

Definition 7: Covariance Graph [6]: The covariance graph $G(\mathbf{\Sigma})$ of $\mathbf{\Sigma}$ is the graph with node set $\mathcal{V} \cup \{Y\}$ and edges between any pair of variables X_i and X_j (or Y and X_i) with $\Sigma_{X_i X_j} > 0$.

Using this definition, we can impose certain structural constraints on the covariance matrix. In particular, we can constrain its corresponding graph to be a tree (i.e., a graph without simple cycles). The following lemma shows that, in this case, we can find an optimal attack in polynomial time.

Lemma 1: If the covariance graph $G(\mathbf{\Sigma})$ is a tree, then removing K of the variables connected to Y in a greedy manner is an optimal attack.

The proof of Lemma 1 can be found in [16].

Building on the above lemma, we can show that the resilient selection problem can also be solved in polynomial time for this special case.

Theorem 3: If the covariance graph $G(\mathbf{\Sigma})$ is a tree, then an optimal resilient observation selection can be found in polynomial time.

The proof of Theorem 3 can be found in [16].

V. Experiments

Having proposed algorithms for both the denial-of-service attack on sensors and the associated resilient sensor (i.e., observation) selection problems, we now proceed to evaluate them experimentally. We aspire to answer three questions:

1) Is the greedy attack algorithm effective in practice? We address this in Section V-B.
2) Is the greedy resilient observation selection algorithm nearly optimal in practice? We address this question in Section V-C.
3) Does the resilient observation selection algorithm significantly improve resilience compared to baseline, and does it lose much compared to baseline under normal (no attack) conditions? We address this in Section V-C.

We focus on the performance of the outputs in terms of posterior variance, and omit detailed results on the running times of the algorithms. Note that, for the dataset used in the experiments, exhaustive search for the optimal solutions was computationally infeasible even for moderately difficult problems (e.g., $K > 10$), while the running times of the greedy algorithms were less than a minute. However, for the problems used for illustration in this section, we were able to find the optimal solutions using exhaustive search.

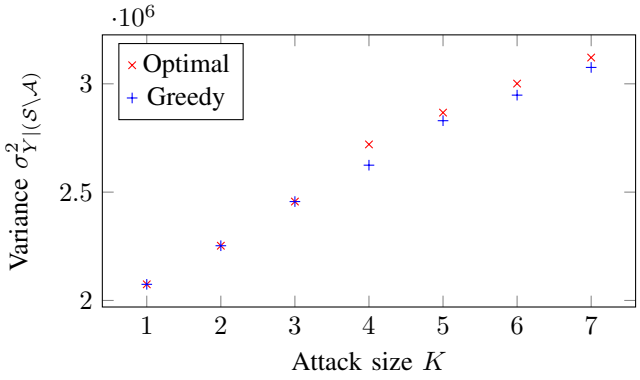

Fig. 1. Posterior variances for optimal and greedy attacks.

A. Datasets

We have applied our algorithms to multiple datasets (e.g., measurements from air quality monitoring stations in Beijing and Shanghai[5], see [22]); however, due to lack of space, we will present results only for one particular dataset, noting that results were consistent throughout all the datasets. This particular dataset is from the Caltrans Performance Measurement System (PeMS), which collects traffic information from sensors spanning the freeway system across all major metropolitan areas of California.[6] We chose 37 stations from the Bay Area with complete hourly data available for January to be the set of possible observations, with each possible observation measuring the traffic flow over a corresponding road segment. For the predictor variable, we used the average over all traffic flow measurements from around the Bay Area (i.e., average traffic situation).

B. Greedy Algorithm for Attack

Figure 1 shows the posterior variance of the predictor variable as a function of the attack size K for optimal and greedy attacks. We can see that the variances resulting from the optimal and greedy attacks are identical in many cases, and there is little difference between the two in the remaining ones. More specifically, the difference is greatest for $K = 4$, but even in this case, it is less than 4%.

C. Greedy Algorithm for Resilient Observation Selection

Figure 2 shows the posterior variance of the predictor variable as a function of the attacks size K for optimal resilient selection (x), greedy resilient selection (o), and optimal non-resilient selection (+) with and without an attack (blue and red, respectively) for $N = 7$ (note that attacks are found using exhaustive search in all cases). Firstly, we can see that the output of the proposed greedy algorithm and the optimal resilient selection are very close, and in many cases, they are identical. Secondly, we can see that the resilient selection performs much better than the non-resilient selection in case of an attack: for small-sized attacks

[5] http://research.microsoft.com/pubs/193973/Air%20Quality%20Data.zip

[6] http://pems.dot.ca.gov/

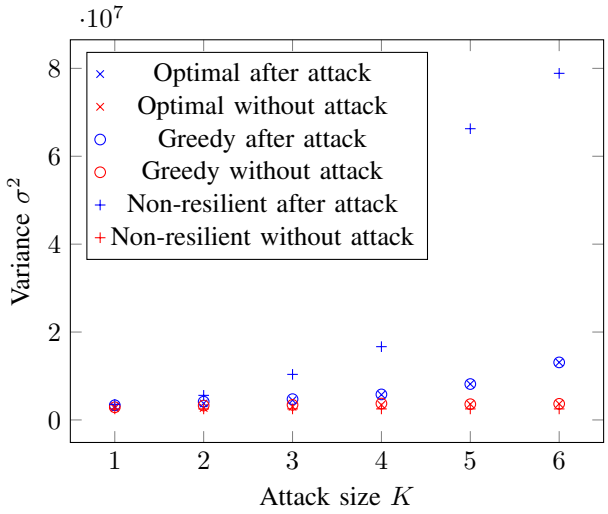

Fig. 2. Posterior variances for optimal and greedy resilient and optimal non-resilient selections with $N = 7$.

($K = 1, 2$), the posterior variance for the non-resilient selection is 17 - 49% higher than for the resilient selection, and for larger attacks, there can be an order of magnitude difference (e.g., 711% and 502% for $K = 5$ and 6). On the other hand, there are much smaller differences between the quality of the two selections (resilient and baseline) when there is no attack: for $K = 1, 2$, the posterior variance for the non-resilient selection is less than 9% lower than for the resilient selection, and the difference remains less than 32% even for higher values of K.

VI. Conclusion

In this paper, we considered the problem of selecting sensor locations for monitoring a large area in an adversarial setting. We formulated the problem of attack-resilient observation selection by assuming Gaussian process regression and worst-case DoS attacks against the sensors. First, we showed that both finding an optimal attack against a given selection and finding an optimal resilient selection are NP-hard problems. Then, we proposed efficient heuristic algorithms for both problems, and established approximation bound on the heuristic algorithm for resilient selection. Next, we considered a special case of the problem, in which the covariance matrix of the variables has a tree structure, and showed that both problems can be solved in polynomial time for this special case. Using experiments based on real-world covariance data, we demonstrated that both of the proposed heuristic algorithms perform exceptionally well in practice. Finally, these numerical results also showed that the increase in posterior variance under normal operations (i.e., without an attack) is quite small, while the decrease in posterior variance in the case of an attack is substantial.

Acknowledgments: This work was supported in part by the National Science Foundation (CNS-1238959), Air Force Research Laboratory (FA8750-14-2-0180), and National Institute of Standards and Technology (70NANB13H169).

References

[1] Y. Bazi, N. Alajlan, and F. Melgani. Improved estimation of water chlorophyll concentration with semisupervised Gaussian process regression. *IEEE Transactions on Geoscience and Remote Sensing*, 50(7):2733–2743, 2012.

[2] J. Blanchet, A. C. Davison, et al. Spatial modeling of extreme snow depth. *Annals of Applied Statistics*, 5(3):1699–1725, 2011.

[3] J. L. Bredin, E. D. Demaine, M. Hajiaghayi, and D. Rus. Deploying sensor networks with guaranteed capacity and fault tolerance. In *Proc. of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 309–319. ACM, 2005.

[4] R. D. Carr, H. J. Greenberg, W. E. Hart, G. Konjevod, E. Lauer, H. Lin, T. Morrison, and C. A. Phillips. Robust optimization of contaminant sensor placement for community water systems. *Mathematical Programming*, 107(1-2):337–356, 2006.

[5] J. Chen, K. H. Low, C. K.-Y. Tan, A. Oran, P. Jaillet, J. M. Dolan, and G. S. Sukhatme. Decentralized data fusion and active sensing with mobile sensors for modeling and predicting spatiotemporal traffic phenomena. In *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 163–173, 2012.

[6] A. Das and D. Kempe. Algorithms for subset selection in linear regression. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 45–54. ACM, 2008.

[7] A. Das and D. Kempe. Submodular meets spectral: Greedy algorithms for subset selection, sparse approximation and dictionary selection. In *Proceedings of the 28th International Conference on Machine Learning (ICML)*, pages 1057–1064, 2011.

[8] G. Davis, S. Mallat, and M. Avellaneda. Adaptive greedy approximations. *Constructive Approximation*, 13(1):57–98, 1997.

[9] S. S. Dhillon and K. Chakrabarty. Sensor placement for effective coverage and surveillance in distributed sensor networks. In *Proceedings of the 2003 IEEE Wireless Communications and Networking Conference (WCNC)*, volume 3, pages 1609–1614. IEEE, 2003.

[10] A. Hahn, A. Ashok, S. Sridhar, and M. Govindarasu. Cyber-physical security testbeds: architecture, application, and evaluation for smart grid. *IEEE Transactions on Smart Grid*, 4(2):847–855, 2013.

[11] M. Ishizuka and M. Aida. Performance study of node placement in sensor networks. In *Proceedings of the 24th International Conference on Distributed Computing Systems Workshops (ICDCS)*, pages 598–603. IEEE, 2004.

[12] P. Kou, F. Gao, and X. Guan. Sparse online warped Gaussian process for wind power probabilistic forecasting. *Applied Energy*, 108:410–428, 2013.

[13] A. Krause, B. McMahan, C. Guestrin, and A. Gupta. Selecting observations against adversarial objectives. In *Proceedings of the 21st Annual Conference on Advances in Neural Information Processing Systems (NIPS)*, pages 777–784, 2007.

[14] A. Krause, B. McMahan, C. Guestrin, and A. Gupta. Robust submodular observation selection. *Journal of Machine Learning Research*, 9:2761–2801, 2008.

[15] A. Krause, A. Singh, and C. Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *Journal of Machine Learning Research*, 9:235–284, 2008.

[16] A. Laszka, Y. Vorobeychik, and X. Koutsoukos. Resilient observation selection in adversarial settings – Extended version with proofs. `http://aronlaszka.com/papers/laszka2015resilient.pdf`, 2015.

[17] I. Nevat, G. W. Peters, and I. B. Collings. Location-aware cooperative spectrum sensing via gaussian processes. In *Proc. of the 13th Australian Comm. Theory Work. (AusCTW)*, pages 19–24. IEEE, 2012.

[18] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

[19] S. Sridhar and G. Manimaran. Data integrity attacks and their impacts on SCADA control system. In *IEEE Power and Energy Society General Meeting*, 2010.

[20] Y. Xie, K. Zhao, Y. Sun, and D. Chen. Gaussian processes for short-term traffic volume forecasting. *Transportation Research Record: Journal of the Transportation Research Board*, 2165(1):69–78, 2010.

[21] M. Younis and K. Akkaya. Strategies and techniques for node placement in wireless sensor networks: A survey. *Ad Hoc Networks*, 6(4):621–655, 2008.

[22] Y. Zheng, F. Liu, and H.-P. Hsieh. U-Air: When urban air quality inference meets big data. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1436–1444. ACM, 2013.