An Analysis of the Mutual Information Upper Bound for Sensor-Subset Selection

Idyano Leroy
SAMOVAR
Télécom SudParis, Institut Polytechnique de Paris
Paris, France
idyano.leroy@telecom-sudparis.eu

Augustin A. Saucan
Institute of Telecommunications
TU Wien
Vienna, Austria
augustin.saucan@tuwien.ac.at

Yohan Petetin
SAMOVAR
Télécom SudParis, Institut Polytechnique de Paris
Paris, France
yohan.petetin@telecom-sudparis.eu

Daniel Clark
VLC ECS
University of Southampton
Southampton, UK
daniel.clark@soton.ac.uk

***Abstract*—The ability to rapidly select an optimal subset of sensors is of critical importance in massive multi-sensor target tracking. Various information metrics exist for selecting the subset of sensors that is most informative with respect to the target being tracked. Moreover, information bounds were proposed as approximate metrics in order to speed up the selection algorithms. In this paper, we provide an analysis on the information loss and its impact on the subset selection problem when employing an information upper bound instead of the exact mutual information metric. We design several greedy sensor-selection algorithms that sequentially evaluate the exact mutual information between a set of sensors and the target. Subsequently, we compare these algorithms with a sensor-selection method that employs an information upper bound and highlight situations where the latter finds sub-optimal solutions.**

***Index Terms*—mutual information, information upper bound, sensor selection, Kalman filter, submodularity.**

I. Introduction

Technological advances in sensor systems have increased the controllability and adaptability of sensing devices, with many commands being controlled autonomously. Indeed, many developments in networked systems and autonomous vehicles have led to configurable networks of sensor systems. More and more new types of sensors are emerging, with device technologies and system architectures that lend themselves to adaptive operation, such as the Internet of Things. This need for decision making is primarily driven by resource constraints, such as energy, bandwidth and processing power. Therefore, decision systems are needed to optimise resources, i.e., to maximise a specific performance criterion under given resource constraints.

In our work, the choice of the performance criterion, or utility function, is driven by:

- The need to quantify the precision of the posterior distribution for a specified action. This is in contrast with methods that only consider second-order moments, e.g., variance.
- The need for a measure that is agnostic of the filtering algorithm used to compute the posterior distribution.

Given these specifications, information theory has emerged as a suitable framework [3]. Indeed methods based on Fisher information, entropy or mutual information are good tools that address both needs.

Sensor management has been the subject of a number of approaches based on information theory, see e.g., [1], [2]. In [4], the authors have stated the general problem of sensor management within the stochastic control framework. [5] explored a number of information-theoretic approaches to sensor management. These include various ways of modelling information gain using statistical divergences. They also provided links to risk and classification. [6] compared different divergence measures for sensor management applied to target tracking. [7] extended [4] for the multiple target tracking problem and used Rényi divergence. [8], [9] proposed information-theoretic approaches applied to multiple sensor control. [10], [11] proposed sensor selection methodology for multi-target multi-sensor tracking and ocean-of-things. More recently, [12] used tools from network coding theory to compute the mutual information between random sets of targets and sensors.

In this paper, we focus on the problem of a single target observed by multiple sensors. We propose an exact and efficient scheme for computing the mutual information between a target and a set of sensors for linear Gaussian models. We provide greedy algorithms that optimise our mutual information metric and compare their efficiency to an exact maximisation of a mutual information upper bound.

The paper is organised as follows. In the next section we describe the technical problem; we present our approach and highlight its difference from related work. In Sec. III, we

This work was supported by the IP Paris *Centre Interdisciplinaire d'Etudes pour la Défense et la Sécurité* (CIEDS) under the NITSM-MTS project with award no. 2022650059 and in part by AFOSR grant FA8655-22-1-7013.

derive the sequential computation of our objective function (2) and give its analytical expression in two specific cases. In Sec. IV, we exploit the properties of (2) to design an algorithm tailored for maximising the proposed information metric. Finally, Sec. V describes simulation results that demonstrate the usefulness of our method.

II. Problem formulation

Let $(X_k)_{k\geq 0}$ be a sequence of hidden random variables, where $X_i \in \mathbb{R}^{d_x}$; at each time step, S sensors produce measurements $(Y_{1,k}, \cdots, Y_{S,k})$, $Y_{s,k} \in \mathbb{R}^{d_{y_s}}$. The general filtering problem is to estimate, at each time step, the posterior distribution of X_k given the past observations. However, in power-constrained applications, activating all sensors is impractical in reality, mainly due to the resources required. Furthermore, for large sensor networks, the small gain obtained by incorporating data from additional sensors can be overshadowed by the computational overhead of processing them. Consequently, at each time step k, the objective is twofold:

(i) select a n_k-size subset of sensors and collect their measurements $\mathbf{y}_{\mathbf{n}_k} = (y_{i_1}, \cdots, y_{n_k})$, where $\mathbf{n}_k$ describes the indices of selected sensors, i.e., $\mathbf{n}_k = \{i_1, \cdots, i_{n_k}\}$;

(ii) update the posterior distribution $p(x_k|\mathbf{y}_{\mathbf{n}_{0:k}})$, where $\mathbf{y}_{\mathbf{n}_{0:k}} = (\mathbf{y}_{\mathbf{n}_0}, \cdots, \mathbf{y}_{\mathbf{n}_k})$, from $p(x_{k-1}|\mathbf{y}_{\mathbf{n}_{0:k-1}})$. In hidden Markov models, this problem is well-known and is solved by the Kalman filter in the linear and Gaussian case [15], or by approximate methods, such as, the Extended Kalman Filter (EKF) or sequential Monte Carlo methods [16], [17].

In this paper, we focus on problem (i), which consists in selecting the best subsets $\{\mathbf{n}_k\}_{k\geq 0}$ according to a performance criterion. To that end, we focus on mutual information. In this context, it quantifies the amount of information that a set of observations provides about the hidden variable. More precisely, the mutual information between X_k and a subset of observations $\mathbf{Y}_{\mathbf{n}_k}$, conditionally on the past $\mathbf{y}_{0:\mathbf{n}_{k-1}}$, is denoted as $\mathcal{I}_{k|k-1}(X_k; \mathbf{Y}_{\mathbf{n}_k})$ and is defined as

$$\mathcal{I}_{k|k-1}(X_k; \mathbf{Y}_{\mathbf{n}_k}) = \mathbb{E}\left(\log\left(\frac{p(X_k, \mathbf{Y}_{\mathbf{n}_k}|\mathbf{y}_{\mathbf{n}_{0:k-1}})}{p(X_k|\mathbf{y}_{\mathbf{n}_{0:k-1}})p(\mathbf{Y}_{\mathbf{n}_k}|\mathbf{y}_{\mathbf{n}_{0:k-1}})}\right)\middle|\mathbf{y}_{\mathbf{n}_{0:k-1}}\right). \quad (1)$$

Note that (1) is an expectation with respect to the pair $(X_k, \mathbf{Y}_{\mathbf{n}_k})$ and coincides with the Kullback-Leibler divergence between the joint distribution $p(x_k, \mathbf{y}_{\mathbf{n}_k}|\mathbf{y}_{\mathbf{n}_{0:k-1}})$ in the original model and that in a model where the considered subset of observations at time k would be conditionally independent of X_k. Consequently, (1) tends to be close to zero when the two distributions are close, i.e., when the observations do not impact the predicted distribution, so $p(x_k|\mathbf{y}_{\mathbf{n}_{0:k}}) \approx p(x_k|\mathbf{y}_{\mathbf{n}_{0:k-1}})$. In order to maximise (1) w.r.t. $\mathbf{n}_k$, we introduce the network information function

$$v_k : \begin{cases} \mathbf{n}_k \in \mathcal{P}(A) \setminus \{\emptyset\} \mapsto \mathcal{I}_{k|k-1}(X_k; Y_{\mathbf{n}_k}) \\ \emptyset \mapsto 0 \end{cases} \quad (2)$$

where $\mathcal{P}(A)$ is the power set of $A = \{1, 2, \cdots, S\}$. Finally, problem (i) is equivalent to maximising v_k, for all $k \in \mathbb{N}$, under the p-cardinality constraint, i.e,

$$\begin{aligned} &\underset{\mathbf{n}_k \in \mathcal{P}(A)}{\text{maximise}} \quad v_k(\mathbf{n}_k) \\ &\text{subject to} \quad |\mathbf{n}_k| \leq p \end{aligned} \quad (\mathrm{P}_k)$$

where $|\mathbf{n}_k|$ denotes the cardinality of $\mathbf{n}_k$. Mutual information driven subset selection has been studied previously in [18], [19]. However, they did not provide an efficient computational scheme. In fact, [18] only studied the performance guarantees of different optimisation schemes, but did not provide an efficient means of computing mutual information. [19] provided a sampling algorithm to compute the conditional entropy of a candidate subset, but limited its application to fully specified graphical models.

In [13] an upper bound of (2) was derived and optimised. Its advantage is that it can be exactly maximised in S function evaluations. However, using this upper bound as an objective function can lead to some problems, as it does not take into account the complementary information between the selected sensors, as we will show later. In addition, the proposed optimisation scheme may not be optimal due to the theoretical properties of such an objective function. Finally, [14] uses a Monte Carlo estimator of the mutual information, but the complexity of their computational scheme grows exponentially with the number of sensors.

We propose a computationally inexpensive method that is able to take into account complementary information from different sensors and, as a result, dynamically select the most informative subset of sensors more accurately. Rather than optimising an upper bound of mutual information, our approach amounts to suboptimally maximise an exact computation of mutual information. To make it practical, it is necessary to derive a method to compute $v_k(\mathbf{n}_k)$ with linear complexity with respect to n_k and integrate it into an optimisation scheme. To this end, we first show that, at a given time k, our objective function $v_k(\mathbf{n}_k)$ can be computed sequentially from any set of cardinality $n_k - 1$ included in $\mathbf{n}_k$. In particular, we give its analytic expression in the linear Gaussian case and its approximated version in the non-linear Gaussian case. We then provide sequential optimisation algorithms that take advantage of this sequential computation and the properties of our metric. In Sec. V, we demonstrate the approach in simulations on different scenarios.

III. Mutual information for sequential sensor subset selection

In this section, we show that $v_k(\mathbf{n}'_k)$ can be computed from $v_k(\mathbf{n}_k)$, where $\mathbf{n}_k$ is a subset of sensors of cardinality $n_k - 1$ included in $\mathbf{n}'_k$. Without loss of generality, in this section we consider the static case; thus, we drop the time index k and focus on computing the mutual information between a r.v. X and a set of observations $\mathbf{Y}_n = (Y_1, \cdots, Y_n)$ from that between X and $\mathbf{Y}_{n-1}$, where $p(x, \mathbf{y}_n) = p(x) \prod_{i=1}^{n} g_i(y_i|x)$.

Starting from

$$\mathcal{I}(X;\mathbf{Y}_n) = \mathbb{E}\left(\log\left(\frac{p(X,\mathbf{Y}_n)}{p(X)p(\mathbf{Y}_n)}\right)\right),$$

it can be deduced from the factorization of $p(x,\mathbf{y}_n)$ that

$$\mathcal{I}(X;\mathbf{Y}_n) = \mathcal{I}(X;\mathbf{Y}_{n-1}) + \mathcal{I}(X;Y_n) - \mathcal{I}(Y_n;\mathbf{Y}_{n-1}). \quad (3)$$

In other words, the computation of $\mathcal{I}(X;\mathbf{Y}_n)$ relies on that of $\mathcal{I}(X;Y_n)$, which can be computed at a fixed computational cost, and that of $\mathcal{I}(Y_n;\mathbf{Y}_{n-1})$, which involves an n-fold integral that is intractable for large sensor networks. Ignoring the latter term leads to the the Mutual Information Upper Bound (MIUB), defined as

$$\mathrm{MIUB}(X,\mathbf{Y}_n) = \sum_{i=1}^{n}\mathcal{I}(X;Y_i) \quad (4)$$

which has been maximised in [13], [14]. This approximation allows for a fast, i.e., linear-complexity with the number of sensors, evaluation of a sensor subset. However, we point out that ignoring $\mathcal{I}(Y_n;\mathbf{Y}_{n-1})$ can lead to significantly sub-optimal strategies, both from a theoretical and an experimental point of view. Indeed, it has the disadvantage of not taking into account the complementary information that different sensors can provide.

A. Mutual information for linear Gaussian sensors

Here we present an exact computation of the network information in the linear Gaussian case. The prior and the likelihood satisfy $p(x) = \mathcal{N}(x;\mu_x,\Sigma_x)$, $g_n(y_n|x) = \mathcal{N}(y_n;H_nx,\Sigma_n)$, where $\mathcal{N}(x;\mu_x,\Sigma_x)$ is the probability density function of the Gaussian distribution with mean μ_x and covariance matrix Σ_x taken at point x. In this particular case, all the distributions $\{p(y_l|\mathbf{y}_{l-1})\}_{2\le l\le n}$ and $\{p(x|\mathbf{y}_l)\}_{1\le l\le n}$ are Gaussian and their densities can be derived from the above assumptions and classical results on Gaussian distributions [20]. In particular, the key observation is that $\mathbb{E}(\log(p(Y_n|\mathbf{y}_{n-1}))|\mathbf{y}_{n-1})$ does not depend on $\mathbf{y}_{n-1}$, so we have the following result.

Proposition III.1. *Assuming that all sensors are described by the linear Gaussian model above, and starting from a subset $\mathbf{n} = \{i_1, i_2, \cdots, i_{n-1}\}$ with an arbitrary ordering of elements and corresponding network information $v(\mathbf{n})$, the network information value of the subset $\mathbf{n}' = \mathbf{n}\cup\{i_n\}$, $\forall i_n \in A\setminus\mathbf{n}$, is obtained as*

$$v(\mathbf{n}') = v(\mathbf{n}) + \frac{1}{2}\log\left(\frac{\det\left(\Sigma_{i_n} + H_{i_n}\widetilde{P}_nH_{i_n}^\top\right)}{\det\left(\Sigma_{i_n}\right)}\right), \quad (5)$$

where $\widetilde{P}_n$ satisfy

$$\begin{cases} \widetilde{P}_n = (I - K_nH_{i_n})\widetilde{P}_{n-1}, & \forall n > 0, \\ \widetilde{P}_0 = \Sigma_x\,, \end{cases} \quad (6)$$

and $K_n = \widetilde{P}_{n-1}H_{i_n}^\top\left(\Sigma_{i_n} + H_{i_n}\widetilde{P}_{n-1}H_{i_n}^\top\right)^{-1}$, for $n > 0$.

Sketch of proof. Starting from (3), we just need to show that $\mathcal{I}(X;Y_n) - \mathcal{I}(Y_n;\mathbf{Y}_{n-1}) = \mathbb{E}\left(\log\left(p(Y_n|X)\right)\right) - \mathbb{E}\left(\log\left(p(Y_n|\mathbf{Y}_{n-1})\right)\right)$ coincides with the incremental term in (5). Let us note that $p(y_n|\mathbf{y}_{n-1})$ can be sequentially propagated from

$$p(y_n|\mathbf{y}_{n-1}) = \int p(y_n|x)\,p(x|\mathbf{y}_{n-1})\mathrm{d}x,$$
$$p(x|\mathbf{y}_{n-1}) = \frac{\prod_{i=1}^{n-1}p(y_i|x)p(x)}{\int\prod_{i=1}^{n-1}p(y_i|x)p(x)\mathrm{d}x}.$$

Next, classical results on Gaussian distributions [20] ensure that

$$p(x|\mathbf{y}_{n-1}) = \mathcal{N}(x;\widetilde{m}_{n-1},\widetilde{P}_{n-1})$$
$$p(y_n|\mathbf{y}_{n-1}) = \mathcal{N}(y_n;H_n\widetilde{m}_{n-1},\Sigma_n + H_n\widetilde{P}_{n-1}H_n^T)$$

where $\widetilde{m}_0 = \mu_x$, $\widetilde{P}_0 = \Sigma_x$ and recursively for $i = 2,3,\ldots,n$

$$\begin{aligned} \widetilde{m}_{i-1} &= \widetilde{m}_{i-2} + K(y_{i-1} - H_{i-1}\widetilde{m}_{i-2})\,, \\ \widetilde{P}_{i-1} &= (I - KH_{i-1})P_{i-2}\,, \\ K &= \widetilde{P}_{i-2}H_{i-1}^T(\Sigma_{i-1} + H_{i-1}\widetilde{P}_{i-2}H_{i-1}^T)^{-1}. \end{aligned}$$

Consequently,

$$-\mathbb{E}\left(\log\left(p(Y_n|\mathbf{Y}_{n-1})\right)\right) = \frac{1}{2}\log\left(\frac{\det\left(\Sigma_n + H_n\widetilde{P}_nH_n^T\right)}{(2\pi)^{d_n}}\right)$$
$$+\frac{1}{2}\int(y_n-\widetilde{m}_{n-1})(\Sigma_n+H_n\widetilde{P}_{n-1}H_n^T)^{-1}(y_n-\widetilde{m}_{n-1})^T p(\mathrm{d}\mathbf{y}_n).$$

The above integral is taken with respect to $\mathbf{y}_n = (y_1, y_2, \cdots, y_n)$. However, after first integrating with respect to y_n, it follows from Lemma A.1 (in the Appendix) that the result does not depend on $\mathbf{y}_{n-1}$. Second, after integrating this result with respect to the remaining variables $\mathbf{y}_{n-1}$, we have

$$-\mathbb{E}\left(\log\left(p(Y_n|\mathbf{Y}_{n-1})\right)\right) = \frac{1}{2}\log\left(\frac{\det\left(\Sigma_n + H_n\widetilde{P}_nH_n^T\right)}{(2\pi)^{d_n}}\right) + \frac{d_n}{2}.$$

Similarly, we get:

$$\mathbb{E}\left(\log\left(p(Y_n|X)\right)\right) = \frac{1}{2}\log\left(\frac{(2\pi)^{d_n}}{\det\left(\Sigma_n\right)}\right) - \frac{d_n}{2},$$

whence (5). □

Note, that the specific forms of the equations (5) and (6) are amenable to a linear complexity implementation with respect to n. In particular, by storing the intermediate results of $v(\mathbf{n})$ and $\widetilde{P}_{n-1}$, we can evaluate $v(\mathbf{n}')$ for all $n > 0$ by performing (6) once and adding the respective information gain to $v(\mathbf{n})$. This fact is exploited by our greedy sensor selection methods presented in Sec. IV.

B. Mutual information for non-linear sensors

We now see the non-linear measurement model. In this case the likelihood becomes $g_n(y_n|x) = \mathcal{N}(y_n; h_n(x), \Sigma_n)$, where h_n is a non-linear function (e.g., in range or bearing measurements), the idea is to linearise h_n around μ as in EKF-based approaches. To that end, the matrix H_n in Prop. III.1 is replaced by

$$H_n \leftarrow \left.\frac{\partial h_n}{\partial x}\right|_{x=\mu_x}$$

where the derivative is evaluated at μ_x.

Note that more sophisticated approximations could be developed, but this is beyond the scope of this paper and our focus is more on the theoretical properties of our network information function.

IV. Sequential network information maximisation

In this section, we present the model of our optimisation framework. The basis of our approach is to sequentially add relevant sensors to build our subset solution. More precisely, if we denote by $\hat{S}_k$ our solution to the problem (P_k), we seek a sequence $(\hat{a}_n)_{1\leq n\leq n_k}$, $\hat{a}_n \in A = \{1, \cdots, S\}$, such that $\hat{S}_k = \bigcup_{n=1}^{n_k} \{\hat{a}_n\}$. In this framework, an optimisation scheme is specified by a particular choice of $(\hat{a}_n)$. The advantage of our fixed cost sequential computation is that the cost of evaluating v_k in (2) remains the same as we add sensors to $\hat{S}_k$.

Problem (P_k) is NP-complete [21], [22]. Although any proposed solution to an NP-complete problem can be verified in polynomial time, there is no known algorithm capable of finding an optimal solution in polynomial time [21]. However, for a monotone and submodular objective function, greedy algorithms have a good approximation ratio [19], [23], [24]. In Sec. III-A we discuss two key properties, monotonicity and submodularity, that ensure the behaviour of an optimisation scheme associated with our network information function. Next, in Sec. III-B we discuss the optimisation schemes, which consist of two greedy maximisation algorithms.

A. Monotonicity and submodularity of network information

We first introduce the monotonicity property.

Proposition IV.1. *The network information function v_k in* (2) *is monotone, i.e., for all $(U, V) \in \mathcal{P}(A) \times \mathcal{P}(A)$ such that $U \subset V$, $v_k(U) \leq v_k(V)$.*

A proof of this result can be found in [19]. This property is particularly useful when we are looking for the best subsets of cardinality at most equal to p, since it guarantees that the best subset is necessarily of cardinality p. Hence, we now set $n_k = p$, for all $k \in \mathbb{N}$.

A second property is the submodularity and its proof can be found in [19].

Proposition IV.2. *The network information function v_k in* (2) *is submodular, i.e., for all $(U, V) \in \mathcal{P}(A) \times \mathcal{P}(A)$ such that $U \subset V$ and for all a such that $a \in A \setminus V$, $v_k(U \cup \{a\}) - v_k(U) \geq v_k(V \cup \{a\}) - v_k(V)$.*

The network information function (2), as a reward function, can be interpreted as the function that maps some benefit to a particular choice of subset of sensors. The submodularity property then ensures that adding sensors has a decreasing marginal benefit. Finally, this property will lead us to propose a suboptimal selection strategy, as we will see later.

B. Maximisation algorithms

We now focus on the problem of maximising v_k in (2). Note that, unlike (4), the computation of $\mathcal{I}(X; \mathbf{Y}_n)$ relies on the quantities computed for that of $\mathcal{I}(X; \mathbf{Y}_{n-1})$. Thus, v_k cannot be maximised directly, and so we discuss two strategies, the *vanilla greedy* and the *randomised greedy* approaches, both of which produce approximate solutions with $\mathcal{O}(Sp)$ function evaluations. We assume that we have already chosen a subset of cardinality $n-1$ and the objective is to include the n-th sensor, for $1 \leq n \leq p$.

1) Vanilla greedy approach: This strategy consist in choosing each $\hat{a}_n$ as the element of $A \setminus \{\hat{a}_1, \cdots, \hat{a}_{n-1}\}$ that maximises the immediate reward,

$$\hat{a}_n = \underset{a \in A \setminus \hat{\mathbf{a}}_{n-1}}{\operatorname{argmax}} \; v_k(\{a\} \cup \hat{\mathbf{a}}_{n-1}) - v_k(\hat{\mathbf{a}}_{n-1}) \tag{7}$$

where $\hat{\mathbf{a}}_{n-1} = \{\hat{a}_1, \cdots, \hat{a}_{n-1}\}$. As shown in [19], [23], [24], for a monotone and submodular function, the previous greedy strategy has a good approximation ratio. More specifically, since v_k is monotone and submodular, the solution returned by the vanilla greedy algorithm achieves a tight $1 - 1/e$ approximation guarantee [19], [23], [24].

2) Randomised greedy: For the second strategy, since in our problem there is no guarantee that a locally optimal choice coincides with a globally optimal choice, we introduce randomness into the previous scheme. We thus consider the immediate reward as a *preference* for a possible choice of $\hat{a}_n$; so at each step, we sample

$$\hat{a}_n \sim P(\hat{a}_n = a)$$

where $P(\hat{a}_n = a)$ is defined as a soft-max distribution

$$P(\hat{a}_n = a) \propto \begin{cases} \exp\left(v_k(\{a\} \cup \hat{\mathbf{a}}_{n-1}) - v_k(\hat{\mathbf{a}}_{n-1})\right), & \text{if } a \in A \setminus \hat{\mathbf{a}} \\ 0, & \text{otherwise} \end{cases},$$

and $\sum_{a \in A} P(\hat{a}_n = a) = 1$.

V. Simulation study

In this section we discuss the previous maximisation algorithms from an experimental point of view. We consider different sensor configurations and the selection of sensors is based on the aforementioned algorithms. In addition, we compare our approach with the direct maximisation of the MIUB:

$$v'_k : \begin{cases} \mathbf{n}_k \in \mathcal{P}(A) \setminus \{\emptyset\} \mapsto \mathrm{MIUB}_{k|k-1}(X_k; Y_{\mathbf{n}_k}) \\ \emptyset \mapsto 0 \end{cases} \tag{8}$$

where $\mathrm{MIUB}_{k|k-1}(X_k;Y_{\mathbf{n}_k}) = \sum_{i\in\mathbf{n}_k} \mathcal{I}_{k|k-1}(X_k;Y_i)$. More specifically, we first consider two static cases to illustrate the advantages and limitations of our solution. We then consider a dynamic scenario to illustrate the benefits of our approach in terms of tracking accuracy.

A. Static linear scenario

In the first simulation study, we consider a linear static scenario. A Gaussian r.v. $X \sim \mathcal{N}(\mu_x, \Sigma_x)$, where

$$\mu_x = \begin{pmatrix} -10 \\ 10 \\ -10 \\ 10 \end{pmatrix}, \; \Sigma_x = \begin{pmatrix} 200 & 0 & 0 & 0 \\ 0 & 200 & 0 & 0 \\ 0 & 0 & 200 & 0 \\ 0 & 0 & 0 & 200 \end{pmatrix}$$

is observed by $S = 200$ sensors. The sensors are separated into two classes: A and B. The first 100 sensors (class A) are described by $g_n(y_n|x) = \mathcal{N}(y_n; H_n x, \Sigma_n)$,

$$H_n = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \; \Sigma_n = \begin{pmatrix} \sigma_i^2 & 0 \\ 0 & \sigma_i^2 \end{pmatrix},$$

where σ_i^2 ranges from 3.20 to 23.0 with 0.2 increments. The other 100 sensors (class B) are described by:

$$H_n = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \; \Sigma_n = \begin{pmatrix} \sigma_i^2 & 0 \\ 0 & \sigma_i^2 \end{pmatrix},$$

where σ_i^2 ranges from 4.0 to 53.5 with 0.5 increments. The 200 sensors are identified by their ID: "A_i" or "B_i" where i ranges from 0 to 99. The optimisation task is to select the subset of cardinality $p = 5$ that maximises v_k.

TABLE I
INFORMATION VALUES (IN NATS) OF THE SOLUTIONS RETURNED BY THE THREE SELECTION ALGORITHMS IN THE FIRST SCENARIO.

Selection algorithm	mean	var	max	min
True maximum value	9.74	.	.	.
Vanilla greedy	9.74	.	.	.
Randomised greedy	7.58	0.41	9.35	4.79
MIUB strategy	5.64	.	.	.

Note that given a sensor class (A or B), the sensors can only be distinguished by the variance of their noise. Since each sensor class provides information about mutually exclusive parts of the target state, we can expect the best subset of sensors to coincide with $\{A_0, B_0, A_1, B_1, A_2\}$. Indeed, an exhaustive search of all possible configurations of size five ensures that this is the subset that maximises mutual information.

We compared three different strategies. The first two are our vanilla greedy and randomised greedy strategies. The third one is the direct maximisation of v'_k in (8), which we call *MIUB strategy*. This maximisation can be achieved exhaustively in S function evaluations, since $\max_{V\in\mathcal{P}(A);|V|\leq p} \mathrm{MIUB}(X;(Y_v)_{v\in V})$ coincides with $\{v_1, \cdots, v_p\}$ whose elements give the p highest values of $\{\mathcal{I}(X;Y_a)\}_{a\in A}$.

The mutual information values corresponding to the solutions found by the three strategies are reported in Table V-A. For the randomised greedy strategy, the results were averaged over 10^4 trials, and we also report the mean, minimum and maximum of the mutual information values in Table V-A.

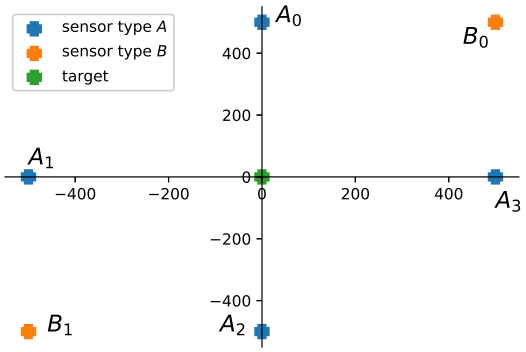

Fig. 1. Sensors and target positions in the second scenario.

We can see that vanilla greedy succeeds in finding the optimal solution in this setting. In fact, it identifies the best sensors, i.e., those with low noise variances, and the right combination of both types of sensors. Randomised greedy gives a slightly worse solution, which is to be expected since it has to perform five random sampling steps. However, most of the subsets returned contain both types of sensors, which is to be expected since it uses the same mutual information calculation scheme as vanilla greedy. The third strategy fails to identify a good sensor synergy. In fact, the mutual information upper bound does not take into account the complementary information provided by different sensors, as highlighted in Sec. III. It identifies the subset $\{A_0, A_1, A_2, A_3, A_4\}$ as maximising the mutual information. This subset is actually made up of the sensors that have the lowest noise variances and are therefore the most informative sensors when considered individually.

TABLE II
INFORMATION VALUES (IN NATS) OF THE SOLUTIONS RETURNED BY THE THREE SELECTION ALGORITHMS IN THE SECOND SCENARIO.

Selection algorithm	mean	max	min
True maximum value	2.20	.	.
Vanilla greedy	2.04	.	.
Randomised greedy	2.02	2.20	1.42
MIUB maximiser	1.52	.	.

B. Static non-linear scenario

We now consider a non-linear static scenario. The distribution of X remains Gaussian with

$$\mu_x = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \; \Sigma_x = \begin{pmatrix} 200 & 0 \\ 0 & 200 \end{pmatrix}$$

and is observed by six range-only sensors, so

$$h_i(x) = \sqrt{(p_x - p_x^i)^2 + (p_y - p_y^i)^2} \tag{9}$$

where p_x and p_y are first and second coordinates of μ_x and (p_x^i, p_y^i) the position of the i-th sensor. Again, we consider two

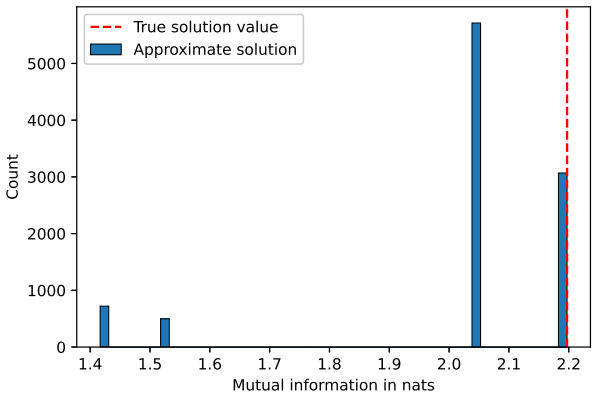

Fig. 2. Histogram of information values returned by the randomised greedy strategy in the second scenario.

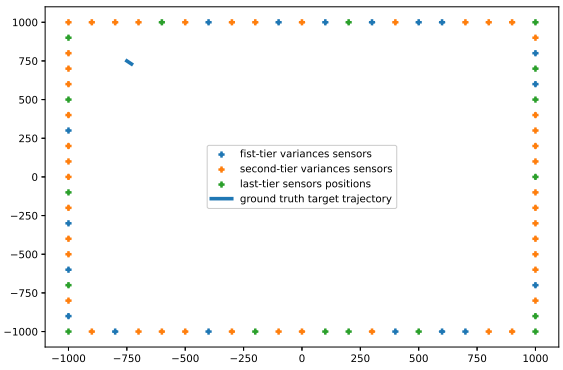

Fig. 3. Sensors and target positions in the third scenario.

classes of sensors, those with a standard deviation of $\sigma_A = 25.0$m (class A) and those with a standard deviation of $\sigma_B = 20.0$m (class B). The first four sensors belong to class A, $\{A_0, A_1, A_2, A_3\}$, and are located at

$$(p_x^0, p_y^0) = (0, 500), \quad (p_x^1, p_y^1) = (-500, 0),$$
$$(p_x^2, p_y^2) = (0, -500), \quad (p_x^3, p_y^3) = (500, 0).$$

Finally, the two class B sensors, $\{B_0, B_1\}$, are positioned at

$$(p_x^4, p_y^4) = (500, 500), \quad (p_x^5, p_y^5) = (-500, -500).$$

The positions of the target and the sensors are given in Fig. 1. The optimisation task is to select the subset of size $p = 2$ that maximises v_k.

In this setting, the parameters that affect the mutual information for a given subset of sensors are their position and their noise variance. In fact, the optimal sensor configuration forms a 90 degree angle with respect to the observed target [25].

As can be seen in Fig. 1, only the pairs $\{A_0, A_1\}, \{A_1, A_2\}, \{A_2, A_3\}, \{A_3, A_0\}$ form an angle of 90 degrees with respect to the target. However, these sensors have the highest noise variances. The optimal pair of sensors is therefore a compromise between well placed sensors and sensors with low noise variance. In fact, due to the symmetry of the experimental setup, there are only four different information values. The pairs of sensors of type A forming an angle of 180 degrees with the target, e.g., $\{A_0, A_2\}$, have the lowest information value $v_k = 1.42$ nats. The pair of sensors of type B forming an angle of 180 degrees with the target, i.e., $\{B_0, B_1\}$, has the second lowest information value $v_k = 1.52$ nats. The pairs of sensors of type A and B, e.g., $\{A_0, B_1\}$, have the second highest information value $v_k = 2.04$ nats. Finally, the pairs of sensors of type A forming an angle of 90 degrees with the target, e.g., $\{A_0, A_1\}$, have the highest information value $v_k = 2.20$ nats.

Again, we compared the three previous strategies. For the randomised greedy strategy, we ran the experiment 10^4 times and computed the average, minimum, and maximum of the mutual information values. The results are reported in Table V-A. We can see that vanilla greedy succeeds in finding the second best information value in this configuration. Since it works sequentially, it first selects sensor B_1 and then searches the remaining sensors to find the one that gives the highest information value, i.e., A_2. However, the MIUB strategy identifies $\{B_0, B_1\}$ as the most informative subset. As in the first study, it fails to identify the good combinations of sensors.

Randomised greedy gives a comparable average information value to the vanilla greedy strategy. However, it succeeds in finding one of the four optimal solutions, $\{A_0, A_1\}, \{A_1, A_2\}, \{A_2, A_3\}, \{A_3, A_0\}$, as shown in the histogram in Fig. 2. The reason for this is that during the first iteration of our maximisation procedure, the algorithm is more likely to choose B_0 or B_1 than one of the A sensors. If it chooses B_0 or B_1, then it is significantly more likely to choose one of the A class sensors than the other B class sensors. Similarly, if it chooses an A class sensor, it is significantly more likely to choose a B class sensor or an A class sensor that forms a 90 degree angle with the target than an A class sensor that forms a 180 degree angle. Thus, if it does not make a locally optimal choice in the first iteration, it will in most cases make an optimal choice in the second iteration. This behaviour is illustrated by the strong asymmetry of the histogram in Fig. 2.

C. Dynamic scenario

We finally consider a non-linear dynamic scenario. At time k, the state vector of a single target is given by $X_k = (p_k^x, p_k^y, v_k^x, v_k^y)^\top$, where p_k^x and p_k^y are the position coordinates and v_k^x and v_k^y are the velocity coordinates. The initial state is $\mu_0 = (-750, 750, 10, -10)^\top$ and we generate a random trajectory from

$$X_{k+1} = FX_k + U_k,$$

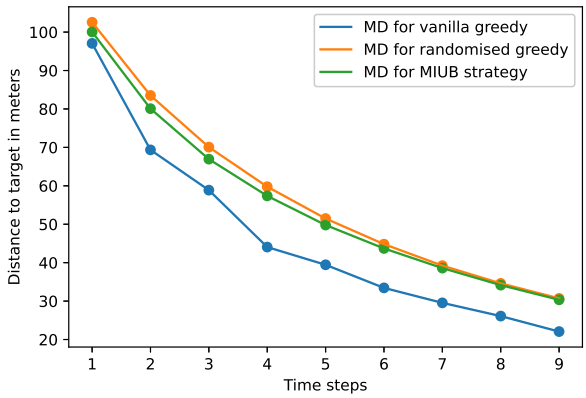

Fig. 4. Mean distance between ground truth target positions and filter estimates for the three strategies in the third scenario.

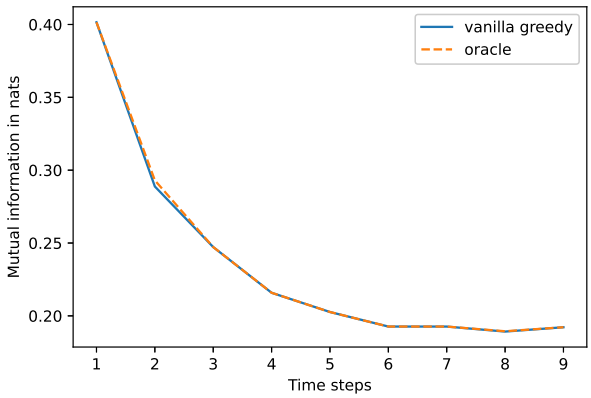

Fig. 5. Average information value of the pair of sensors selected by the vanilla greedy strategy in the third scenario. The average maximum information value is displayed in orange.

where U_k is an independent white Gaussian noise, $U_k \sim \mathcal{N}(0, Q)$ with

$$F = \begin{pmatrix} 1 & 0 & 0.1 & 0 \\ 0 & 1 & 0 & 0.1 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

and

$$Q = \begin{pmatrix} 2.5\,e^{-7} & 0 & 5.0\,e^{-6} & 0 \\ 0 & 2.5\,e^{-7} & 0 & 5.0e^{-6} \\ 5.0\,e^{-6} & 0 & 1.0\,e^{-4} & 0 \\ 0 & 5.0\,e^{-6} & 0 & 1.0\,e^{-4} \end{pmatrix}$$

We have $S = 80$ sensors that can observe the range of the target (see (9)). They are equidistantly placed on a square described by the coordinates $(-1000, 1000)$, $(-1000, -1000)$, $(1000, -1000)$ and $(1000, 1000)$, see Fig. 3. Their noise variances are uniformly sampled in the range $[400, 1000]$ m^2. The goal is to select $p = 2$ sensors at each time step. Once they have been selected, the tracking is ensured by an Extended Kalman Filter.

The previous three selection strategies are used at each time step k to select the most informative sensor pair. To measure the impact of the selection algorithm on the quality of the tracking, we compute the mean distance (MD) between the ground truth target positions and the filter estimates of the corresponding time. The results are shown in Fig. 4. We can see that the vanilla greedy strategy outperforms the MIUB in this setting. This gain is due to our information metric. The MIUB metric overcounts information, it combines too much redundant information and it does not explore the complementarity between sensors. We can see that in this configuration, the randomised strategy gives the worst result, which is to be expected given the very large number of sensors.

In order to discuss the relevance of the solutions obtained with the three strategies, in terms of mutual information value, we compute an *oracle* solution using an exhaustive search. For each strategy and at each time step, the oracle calculates the maximum mutual information value so that we can compare it with the value returned by the selection strategy under consideration. Note that the filtering is not updated with the measurements associated with the oracle, so the oracle depends on the chosen strategy. The results are displayed in Figs. 5, 6, 7. The decreasing trend of the information curve over time for the three strategies is consistent with the fact that tracking reduces uncertainty about the target state. In Fig. 5, the vanilla greedy and oracle curves are almost identical. It shows that vanilla greedy successfully identifies the most informative pair of sensors. Fig.6 showcases the difference between the information curves corresponding to the randomised greedy and the corresponding oracle strategy. On average, the randomised greedy strategy fails to identify the most informative sensor pair. In Fig. 7, we see that both curves are identical for the first iterations of the filter. However, as time passes, the discrepancy between the maximum information value, i.e., the oracle, and the information value of the sensor pair identified by the MIUB strategy increases. This can be explained by the fact that EKF filtering reduces the second order uncertainty of the hidden state over time. Therefore, sensors with low noise variances (which are typically selected by MIUB) are less advantageous in this case. Instead, complementary combinations of sensors (as captured by the network information metric) lead to better target estimates, as shown by the much closer fit of the vanilla greedy algorithm to the oracle in Fig. 5.

VI. Conclusion

This paper explores the use of mutual information to facilitate sensor selection for estimating a target from multiple sensors. Much previous work on information-theoretic sensor control has typically focused on the selection of a single sensor or has used approximate metrics that do not take into account sensor complementarity. In this paper, we propose a network information metric that takes into account sensor complementarity and analyse this metric in several sensor selection scenarios. We highlight regimes where using

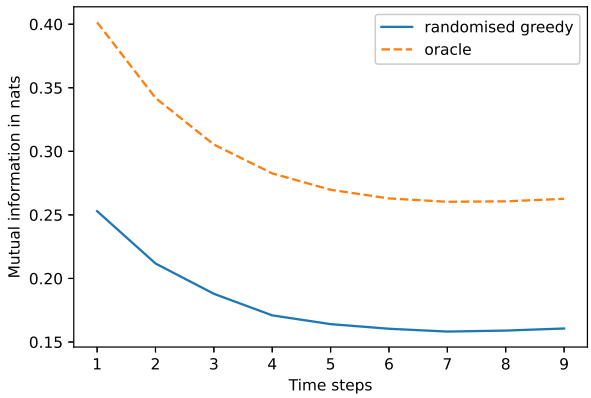

Fig. 6. Average information value of the pair of sensors selected by the randomised greedy strategy in the third scenario. The average maximum information value is displayed in orange.

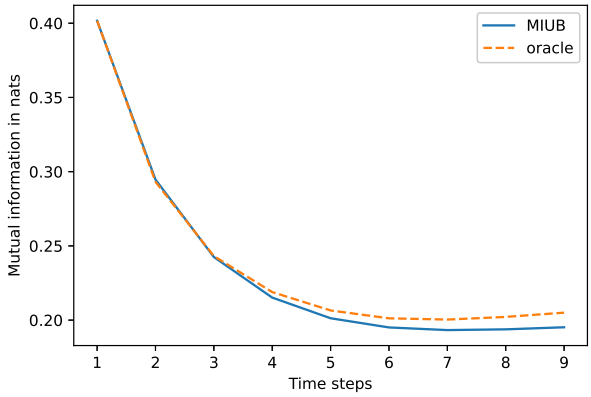

Fig. 7. Average information value of the pair of sensors selected by the MIUB strategy in the third scenario. The average maximum information value is displayed in orange.

the approximate metric leads to suboptimal solutions, while greedy maximisation of our network information metric leads to superior performance.

APPENDIX

In this appendix, we state the lemma used in the proof of Prop. III.1

Lemma A.1. *Let X be a random variable with probability distribution $\mathbb{P}_X$ such that $\mathbb{E}(X) = \mu$ and $\mathbb{V}(X) = P$. Then, for any covariance matrix Q:*

$$\mathbb{E}\left((HX-b)^T Q(HX-b)\right) = \mathrm{Tr}(QHPH^T) + (H\mu - b)^T Q(H\mu - b),$$

where $\mathrm{Tr}(M)$ denotes the trace of matrix M.

REFERENCES

[1] A. O. Hero and D. Cochran, "Sensor Management: Past, Present, and Future," in IEEE Sensors Journal, vol. 11, no. 12, Dec. 2011, pp. 3064-3075.

[2] V. Krishnamurthy, M. Mallick, and B. N. Vo. "Radar Resource Management for Target Tracking-A Stochastic Control Approach," Integrated Tracking, Classification, and Sensor Management: Theory and Applications, Wiley-IEEE Press, 2012, pp. 417-446.

[3] A. O. Hero III, C. M. Kreucher, and D. Blatt, "Information theoretic approaches to sensor management," Foundations and applications of sensor management, 2008, pp. 33-57.

[4] D. A. Castañón and Lawrence Carin, "Stochastic control theory for sensor management," Foundations and applications of sensor management, Boston, MA, Springer US, 2008, pp. 7-32.

[5] C. M. Kreucher, A. O. Hero, K. D. Kastella and M. R. Morelande, "An Information-Based Approach to Sensor Management in Large Dynamic Networks," in Proceedings of the IEEE, vol. 95, no. 5, May 2007, pp. 978-999.

[6] Yang, Chun, et al. "Comparison of information theoretic divergences for sensor management," Signal Processing, Sensor Fusion, and Target Recognition XX. vol. 8050, SPIE, 2011, pp. 199-208.

[7] Branko Ristic, Ba-Ngu Vo, and Daniel Clark, "A note on the reward function for phd filters with sensor control," IEEE Transactions on Aerospace and Electronic Systems, vol. 47 no. 2, 2011, pp. 1521–1529.

[8] Wang, Xiaoying, et al. "Multi-sensor control for multi-object Bayes filters," Signal Processing, vol. 142, 2018, pp. 260-270.

[9] M. Jiang, W. Yi and L. Kong, "Multi-sensor control for multi-target tracking using Cauchy-Schwarz divergence," 2016 19th International Conference on Information Fusion (FUSION), Heidelberg, Germany, 2016, pp. 2059-2066.

[10] A. Saucan, M. J. Coates and Michael Rabbat, "A multisensor multi-Bernoulli filter," IEEE Transactions on Signal Processing, vol. 65, no 20, 2017, pp. 5495-5509.

[11] A. Saucan, and M. Z. Win, "Information-seeking sensor selection for ocean-of-things," IEEE Internet of Things Journal, vol. 7, no 10, 2020, pp. 10072-10088.

[12] D. E. Clark,"Multi-Sensor Network Information for Linear-Gaussian Multi-Target Tracking Systems," in IEEE Transactions on Signal Processing, vol. 69, 2021, pp. 4312-4325.

[13] N. Cao, S. Choi, E. Masazade and P. K. Varshney, "Sensor Selection for Target Tracking in Wireless Sensor Networks With Uncertainty," in IEEE Transactions on Signal Processing, vol. 64, no. 20, 15 Oct.15, 2016, pp. 5191-5204.

[14] G. M. Hoffmann and C. J. Tomlin, "Mobile sensor network control using mutual information methods and particle filters," IEEE Transactions on Automatic Control, vol. 55 no. 1, 2009, pp.32–47.

[15] R. E. Kalman, "A new approach to linear filtering and prediction problems," 1960, pp. 35-45.

[16] B. Ristic, S. Arulampalam, and N. Gordon, "Beyond the Kalman filter: Particle filters for tracking applications," Artech house, 2003.

[17] F. Desbouvries, Y. Petetin, B. Ait-El-Fquih, "Direct, prediction-and smoothing-based Kalman and particle filter algorithms," Signal processing, vol. 91, no 8, 2011, pp. 2064-2077.

[18] J. L. Williams, J. W. Fisher III, and Alan S. Willsky, "Performance guarantees for information theoretic active inference," Artificial Intelligence and Statistics, PMLR, 2007.

[19] A. Krause, and C. Guestrin, "Near-optimal Nonmyopic Value of Information in Graphical Models," in Proc. of Uncertainty in Artificial Intelligence (UAI), 2005.

[20] C.R. Rao, et al. "Linear statistical inference and its applications," Vol. 2, New York: Wiley, 1973.

[21] U. Feige, "A threshold of ln n for approximating set cover," in Journal of the ACM (JACM), vol. 45 no. 4, ACM New York, NY, USA, 1998, pp. 634–652.

[22] Iyer, Rishabh, et al. "Submodular combinatorial information measures with applications in machine learning," Algorithmic Learning Theory. PMLR, 2021.

[23] G. L. Nemhauser, L. A. Wolsey, and M. L. Fisher, "An analysis of approximations for maximizing submodular set functions—I," Mathematical programming, vol. 14, 1978, pp. 265-294.

[24] A. Krause, Andreas, and D. Golovin, "Submodular function maximisation," Tractability, vol. 3, 2014, pp. 71-104.

[25] A.N. Bishop, B. Fidan, B.D. Anderson, K. Doğançay, and P.N. Pathirana, "Optimality analysis of sensor-target localization geometries," Automatica, vol. 46, no 3, 2010, pp. 479-492.