# Performance metric in closed-loop sensor management for stochastic populations

Emmanuel D. Delande, J. Houssineau, and Daniel E. Clark

***Abstract*—Methods for sensor control are crucial for modern surveillance and sensing systems to enable efficient allocation and prioritisation of resources. The framework of partially observed Markov decision processes enables decisions to be made based on data received by the sensors within an information-theoretic context. This work addresses the problem of closed-loop sensor management in a multi-target surveillance context where each target is assumed to move independently of other targets. Analytic expressions of the information gain are obtained, for a class of exact multi-target tracking filters are obtained and based on the Rényi divergence. The proposed method is sufficiently general to address a broad range of sensor management problems through the application-specific reward function defined by the operator.**

## I. Introduction

Effective sensor management requires strategies for controlling networks of sensors to aid decision-making based on the data that can be acquired from them [1]. This is a very general definition since there are many possible controls of individual sensors for different purposes. We adopt an information-theoretic approach based on partially observable Markov decision processes [2], [3]. This paper considers the problem of sensor management in the context of multi-target tracking. This is particularly important in surveillance applications, where the number of targets is unknown and time-varying, and there is uncertainty about the origin of measurements produced by the sensor.

Estimating the global population of targets within the same probabilistic framework enables decisions about sensor control to be taken in a unified way that accounts for the information that can be determined from the different target sources [4]. Recent studies in sensor control for multi-object filtering have demonstrated the effectiveness of this approach using Finite Set Statistics [5]–[9].

Our solution builds on this approach and considers the estimation the global population, though focusses on the information that can be extracted from individual targets. The approach is based on a model that represents a population of targets in a holistic and global way [10]. This model exploits the advantages of representing the global population, which is important for sensor control, whilst maintaining individual target identities. This enhanced capability enables greater flexibility in the choice of reward function, which can focus on individual targets, classes of targets, or the whole population.

E. D. Delande, J. Houssineau, and D. E. Clark are with the School of Engineering & Physical Sciences, Heriot-Watt University, Edinburgh, UK (e-mails: E.D.Delande@hw.ac.uk, jh207@hw.ac.uk and d.e.clark@hw.ac.uk).

## II. Description of the problem

### A. Surveillance activity

The objective of the surveillance activity is to gather information about a population $\mathcal{X}$ of individuals while they lie in some defined region of the physical space called the surveillance scene. The existence of an individual of the population is assumed certain and not inherent to its presence in the scene.

The time is indexed according to some integer variable $t$. At each instant $t \geq 0$, individuals in the scene are represented via a state $\mathbf{x}$, belonging to some target state space $\mathbf{X}_t$, describing physical and measurable characteristics of the individual which are unknown, but of interest to the operator (position, velocity, etc.). The state space is extended with the empty state $\psi$ to $\bar{\mathbf{X}}_t = \{\psi\} \cup \mathbf{X}_t$, in order to allow the description of individuals currently away from the scene.

The aim of the surveillance problem is then to determine, at any time $t \geq 0$ relevant to the surveillance activity, the state of each individual in the augmented state space, i.e. whether the individual is in the scene and, if so, the value of its characteristics of interest.

### B. Sensor system and observation process

While the size of the population $\mathcal{X}$ and the state of the existing individuals are unknown to the operator, the surveillance scene is observed by some sensor system. At each time $t \geq 0$, the sensor may produce at most one observation $z$, in the observation space $\mathbf{Z}_t$, for each individual in $\mathcal{X}$ currently in the scene. Conversely, it is assumed that any target-generated observation originates from a single individual. The collected observations form an observation set $Z_t = \{z_t^1, \ldots, z_t^n\}$, and it is implicitly assumed that all the observations produced by the sensor are distinct. The observation set $Z_t$ may also contain spurious observations, also called *false alarms*.

### C. Input: current information on the target population

Since the population $\mathcal{X}$ is only accessible through an imperfect observation process, the estimation of the number of targets and their current state is uncertain and evolves across time with the availability of new observations. From now on, $t > 0$ designs some current time step where a sensor management decision is scheduled (see Section II-D).

An individual in the population is assumed to enter the scene (or state space) at some time step $t_\bullet$ and possibly leave it at a later time $t_\circ$. The operator identifies a *potential* individual of the population $\mathcal{X}$ – or *target* – through a *track* composed of a sequence of observations allegedly produced by the individual in the previous times. We denote $\bar{Z}_{t'} = \{\phi\} \cup Z_{t'}$

the observation set at time $t'$ augmented with the empty observation $\phi$. Assuming that an individual is always detected at its time of birth $t_\bullet$, a track can be characterised before the update at time $t$ by the *observation path*

$$y = (\phi, \dots, \phi, z_{t_\bullet}, z_{t_\bullet+1}, \dots, z_{t-1}) \tag{1}$$

where $z_{t_\bullet} \in Z_{t_\bullet}$ and $z_{t'} \in \bar{Z}_{t'}$ for any $t_\bullet < t' \leq t-1$. We implicitly assume that individuals with state $\psi$ are never detected by the sensor, i.e., that individuals can only be detected while they lie in the scene. We denote by $Y_{t-1}$ the set of all possible observation paths at time $t$, before the availability of the current observation set $Z_t$ (see Figure 1).

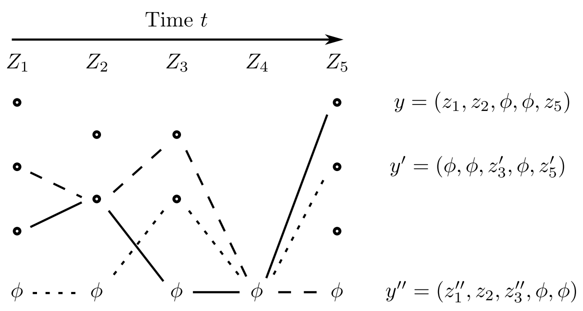


Fig. 1. En example of a few possible tracks at time $t = 6$, before the availability of the observation set $Z_6$. The target characterised by $y'$ entered the scene at time 3 and was detected through the observation $z'_3$, was miss-detected at time 4 and produced the observation $z'_5$ at time 5.

The total information on the population maintained by the operator, at time $t$ and before the availability of $Z_t$, is described by a stochastic population $\mathfrak{Y}_t$ [10]. It is decomposed into two stochastic populations $\mathfrak{Y}^{\mathrm{d}}_t$, $\mathfrak{Y}^{\mathrm{a}}_t$ describing the previously detected and the appearing individuals respectively. A stochastic population can be seen as a process that gives to each possible set of tracks its probability for representing the "true" population $\mathcal{X}$.

*1) Distinguishable individuals:* Each previously detected target is characterised by its observation path and is thus considered as *distinguishable*. An hypothesis $h$ is then defined as a given set of tracks in $Y_{t-1}$ that (allegedly) proposes an accurate representation of the population $\mathcal{X}$. The probability of the hypotheses are described by a distribution $c_t$ on the set of hypotheses $H_t$, which thus satisfies

$$\sum_{h \in H_t} c_t(h) = 1. \tag{2}$$

The distribution $c_t$ can be seen as an "extended cardinality distribution" which characterises not only the number of targets, but also their identity. To each target, represented by some track $y \in Y_{t-1}$, is associated a probability measure $p^y_t \in \mathcal{P}(\bar{\mathbf{X}})$[1], which describes the current state of the target. Thus, each hypothesis $h \in H_t$ can be described by the product measure

$$P^h_t = \bigotimes_{y \in h} p^y_t. \tag{3}$$

Combining (2) and (3), the probability associated to the stochastic population $\mathfrak{Y}^{\mathrm{d}}_t$ is found to be

$$P^{\mathrm{d}}_t = \sum_{h \in H_t} c_t(h) P^h_t. \tag{4}$$

*2) Indistinguishable individuals:* On the other hand, the appearing individuals have not yet been associated to any observation, and are *indistinguishable* in the sense that no specific information is available on any of them. As a consequence, all the appearing targets have the same distribution $p^{\mathrm{a}}_t \in \mathcal{P}(\mathbf{X})$ and their number is driven by a cardinality distribution $c^{\mathrm{a}}_t$. The stochastic population $\mathfrak{Y}^{\mathrm{a}}_t$ is thus characterised by the probability measure

$$P^{\mathrm{a}}_t = \sum_{n \geq 0} c^{\mathrm{a}}_t(n) P^{\mathrm{a},n}_t, \tag{5}$$

where $P^{\mathrm{a},n}_t = (p^{\mathrm{a}}_t)^{\otimes n}$. The distribution $p^{\mathrm{a}}_t$ and the cardinality $c^{\mathrm{a}}_t$ are assumed known to the operator as model parameters. The knowledge of the operator on the current state of the estimated population $\mathfrak{Y}_t$, prior to the collection of the observation set $Z_t$ from the sensor, is thus described by the probability

$$P_t = P^{\mathrm{d}}_t \otimes P^{\mathrm{a}}_t. \tag{6}$$

### D. Sensor management problem

The sensor system producing the observations is imperfect and the uncertainties in the sensing capabilities are described by a stochastic model, assumed known to the operator. We further assume that several *sensor actions* $u \in U_t$ are available at the current time step, each one corresponding to a set of physical parameters (e.g. direction of focus of the sensor, beam width, pulse, etc.) shaping the sensing capabilities and, accordingly, the stochastic model of the sensor.

For each sensor action $u \in U_t$, the observation space of the sensor, assumed finite, is denoted by $\mathbf{Z}_u$. The stochastic model of the sensor under action $u$ – observation noise, probability of detection and probability of false alarm – is described by *a*) a non-negative real-valued function $g_u(z, \cdot) \in \mathcal{B}(\bar{\mathbf{X}}_t)$[2], interpreted as a likelihood and defined for any $z \in \bar{\mathbf{Z}}_u$, where $\bar{\mathbf{Z}}_u = \{\phi\} \cup \mathbf{Z}_u$ is the augmented observation space, and *b*) a probability of false alarm $p_{\mathrm{fa},u}$ on $\mathbf{Z}_u$. For any observation $z \in \mathbf{Z}_u$, it is practical to write $g_u(z, \cdot)$ through a restricted likelihood $\ell_u(z, \cdot) \in \mathcal{B}(\mathbf{X}_t)$ and a probability of detection $p_{\mathrm{d},u}$ on $\bar{\mathbf{X}}_t$ such that:

$$\begin{cases} g_u(z, x) = p_{\mathrm{d},u}(x)\ell_u(z, x), & z \in \mathbf{Z}_u, x \in \mathbf{X}_t, \\ g_u(\phi, x) = 1 - p_{\mathrm{d},u}(x), & x \in \mathbf{X}_t, \\ g_u(z, \psi) = 0, & z \in \mathbf{Z}_u, \\ g_u(\phi, \psi) = 1. \end{cases} \tag{7}$$

In the closed-loop sensor management problem [11], one and only one sensor action $u \in U_t$ must be chosen by the operator in order to drive the observation process that will produce the next observation set $Z_t$. One must design

[1] $\mathcal{P}(\mathbf{X})$ denotes the set of all the probability measures on the measurable space $(\mathbf{X}, \mathbf{B}_\mathbf{X})$, where $\mathbf{B}_\mathbf{X}$ is the Borel $\sigma$-algebra on $\mathbf{X}$.

[2] $\mathcal{B}(\mathbf{X})$ denotes the Banach space of all the bounded and measurable functions on $\mathbf{X}$ equipped with the uniform norm $\|\cdot\|$.

a deterministic policy that, based on the current information on the target population $P_t$ and the expected information $\hat{P}_u$ should the action $u \in U_t$ be chosen, produces a *reward* assessing the *information gain* of each possible action $u$. The sensor is then controlled according to the action that yields the highest expected reward (see Figure 2).

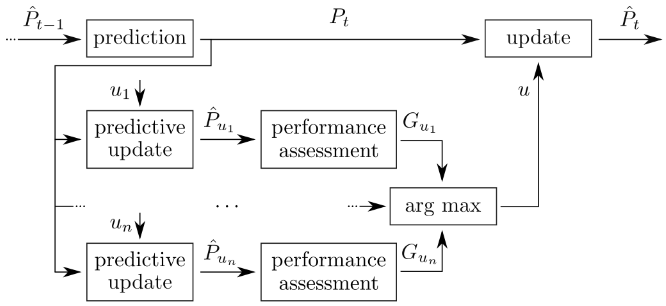


Fig. 2. Principle of close-loop sensor management (time $t$).

## III. Information gain

From now on, $u$ designs an arbitrary choice among the pool of available actions $U_t$, and the rest of the paper focusses on the construction of the information gain $G_u$ (see Figure 2). $Z$ will denote an arbitrary finite subset of $\mathbf{Z}_u$, i.e. a possible observation set collected from the sensor under action $u$, and $\mathfrak{Y}_u$ will denote the corresponding updated stochastic population. The Rényi divergence [12] will be exploited in order to quantify the information gain on each target state.

We will first establish the general form of the information gain $G_u$ through Sections III-A, III-B, and III-C. Emphasizing the information gain in some specific regions of the scene and/or for some specific targets may be valuable depending of the context of the surveillance activity, and we will address this problem in Section III-D.

### A. Individual information gain

The updated set of tracks $\hat{Y}_u$ describing the updated population $\mathfrak{Y}_u$ is composed of the tracks of the form:

- $y:z$[3], where $y \in Y_{t-1}$ and $z \in \bar{Z} = \{\phi\} \cup Z$, i.e. the target characterised by the track $y$ has either generated the observation $z \in Z$ or has been miss-detected ($z = \phi$);
- $\boldsymbol{\phi}_{t-1}:z$[4], where $z \in Z$, i.e. an appearing target has generated the observation $z \in Z$.

The case $\boldsymbol{\phi}_{t-1}:z$ will be denoted $\mathrm{a}:z$ to underline the associated interpretation. We shall now detail the information gain for these two types of tracks.

*1) Updated tracks of the form "$y:z$":* The probability distribution $p_u^{y:z} \in \mathcal{P}(\bar{\mathbf{X}}_t)$ of a new track of the form $y:z$ is given by the Boltzmann-Gibbs transformation [13]:

$$p_u^{y:z}(\mathrm{d}x) = \Psi_{g_u(z,\cdot)}(p_t^y)(\mathrm{d}x) \tag{8a}$$

$$= \frac{1}{p_t^y(g_u(z,\cdot))} g_u(z,x) p_t^y(\mathrm{d}x). \tag{8b}$$

[3]":" is the concatenation operator, i.e. $(e_1, \ldots, e_n):e = (e_1, \ldots, e_n, e)$.

[4]"$\boldsymbol{\phi}_t$" is the sequence of $t$ empty observations, i.e. $\boldsymbol{\phi}_t = (\underbrace{\phi, \ldots, \phi}_{t \text{ times}})$.

We will assume that both $p_t^y$ and $p_u^{y:z}$ admit Radon-Nikodym derivatives with respect to some reference measure $\mu$ on $\bar{\mathbf{X}}_t$ and will denote them the same way for the sake of simplicity. The information gain for an updated track $y:z$ is then defined as the Rényi divergence [12] from $p_t^y$ to $p_u^{y:z}$

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \Big[ \int \big[p_t^y(x)\big]^{\alpha} \big[p_u^{y:z}(x)\big]^{1-\alpha} \mu(\mathrm{d}x) \Big], \tag{9}$$

where $0 < \alpha < 1$ is the order of the divergence[5]. Note that $G_u^{y:z}$ is non-negative, and equals zero if and only if $p_t^y$ coincides with $p_u^{y:z}$ on the augmented state space $\bar{\mathbf{X}}_t$; that is, if the observation $z$ carried no additional information on the target regarding: *a*) its localization in the scene, since $p_t^y = p_u^{y:z}$ on $\mathbf{X}_t$, and *b*) its presence in the scene, since $p_t^y(\psi) = p_u^{y:z}(\psi)$.

*2) New tracks of the form "$\mathrm{a}:z$":* Likewise, the probability distribution $p_u^{\mathrm{a}:z} \in \mathcal{P}(\mathbf{X}_t)$ of a new track of the form $\mathrm{a}:z$ is given by the Boltzmann-Gibbs transformation (8) of the probability distribution of the appearing targets $p_t^{\mathrm{a}}$. Similarly to Section III-A1, we define the information gain $G_u^{\mathrm{a}:z}$ for a new track $\mathrm{a}:z$ as the Rényi divergence from $p_t^{\mathrm{a}}$ to $p_u^{\mathrm{a}:z}$.

### B. Population information gain

We can now determine the information gain assuming a given hypothesis $h \in H_t$ and a given number $n \in \mathbb{N}$ of appearing targets as a representation of the population $\mathcal{X}$. The updated set of tracks $\hat{h}$ is defined as the set of all possible association schemes $\boldsymbol{a} = (h, n, \boldsymbol{h})$, where $\boldsymbol{h}$ represents a specific data association between the tracks in $h$, the $n$ appearing targets and the current observation set $Z$. An association $\boldsymbol{h}$ is an element of the set

$$\begin{aligned} \mathrm{Adm}_Z(h,n) = \Big\{ (h_\mathrm{d}, Z_\mathrm{d}, Z_\mathrm{a}, \nu) \,|\, h_\mathrm{d} \subseteq h,\ Z_\mathrm{d} \subseteq Z_t, \\ Z_\mathrm{a} \subseteq Z_t \setminus Z_\mathrm{d},\ |Z_\mathrm{a}| = n,\ \nu \in \mathcal{S}(h_\mathrm{d}, Z_\mathrm{d}) \Big\}, \end{aligned} \tag{10}$$

where $h_\mathrm{d}$ designs the tracks that are currently detected, $Z_\mathrm{d}$ the observations associated to these detected tracks, $Z_\mathrm{a}$ the observations associated to the $n$ appearing targets, and $\nu$ the bijective function associating detected tracks to observations in $Z_\mathrm{d}$. For some association scheme $\boldsymbol{a} = (h, n, \boldsymbol{h})$, the corresponding updated set of tracks $\hat{h}$ can be characterised by a probability distribution of the form

$$\hat{P}_u^{\boldsymbol{a}} = P_{\mathrm{d},u}^{\boldsymbol{a}} \times P_{\mathrm{md},u}^{\boldsymbol{a}} \times P_{\mathrm{fa},u}^{\boldsymbol{a}}, \tag{11}$$

where

$$P_{\mathrm{d},u}^{\boldsymbol{a}} = \Big[ \prod_{z \in Z_\mathrm{a}} p_t^{\mathrm{a}}(g_u(z,\cdot)) \Big] \Big[ \prod_{y \in h_\mathrm{d}} p_t^y(g_u(\nu(y),\cdot)) \Big],$$

$$P_{\mathrm{md},u}^{\boldsymbol{a}} = \prod_{y \in h \setminus h_\mathrm{d}} p_t^y(g_u(\phi,\cdot)),$$

$$P_{\mathrm{fa},u}^{\boldsymbol{a}} = \Big[ \prod_{z \in Z} 1 - p_{\mathrm{fa},u}(z) \Big] \prod_{z \in Z \setminus (Z_\mathrm{d} \cup Z_\mathrm{a})} \frac{p_{\mathrm{fa},u}(z)}{1 - p_{\mathrm{fa},u}(z)}.$$

[5]The order of the divergence may be dependent on the track $y$, but for the sake of simplicity we will drop this dependency.

The scalar $P^{\boldsymbol{a}}_{\mathrm{d},u}$ (resp. $P^{\boldsymbol{a}}_{\mathrm{md},u}$, $P^{\boldsymbol{a}}_{\mathrm{fa},u}$) represents the probability of the detected individuals (resp. undetected individuals, false alarms) when associated to the observations in $Z$ in the way described by $\boldsymbol{h}$. The updated tracks in $\hat{h}$ are of the form

$$\hat{h} = \Big[\bigcup_{y\in h_{\mathrm{d}}} \{y{:}\nu(y)\}\Big] \cup \Big[\bigcup_{y\in h\setminus h_{\mathrm{d}}} \{y{:}\phi\}\Big] \cup \Big[\bigcup_{z\in Z_{\mathrm{a}}} \{\mathrm{a}{:}z\}\Big]. \quad (12)$$

The information gain from the representation $(h, n)$ to the updated representation $\hat{h}$ can then be expressed in terms of the individual gains (9) as

$$G^{\boldsymbol{a}}_u = \sum_{y\in h_{\mathrm{d}}} G_u^{y:\nu(y)} + \sum_{y\in h\setminus h_{\mathrm{d}}} G_u^{y:\phi} + \sum_{z\in Z_{\mathrm{a}}} G_u^{\mathrm{a}:z}. \quad (13)$$

Considering the updated hypotheses created from all the possible data associations (10), we find the information gain for the pair $(h, n)$, through the observation set $Z$, to be

$$G^{h,n}_u(\cdot|Z) = \sum_{\boldsymbol{h}\in\mathrm{Adm}_Z(h,n)} \hat{P}^{\boldsymbol{a}}_u G^{\boldsymbol{a}}_u. \quad (14)$$

Finally, the information gain for the pair $(h, n)$, expected over all possible observation sets, is found to be

$$G^{h,n}_u = \sum_{Z\subseteq \mathbf{Z}_u} G^{h,n}_u(\cdot|Z). \quad (15)$$

Recalling that the pair $(h, n)$ can be seen as a realisation of the stochastic population $\mathfrak{Y}_t$, it appears that the expected information gain for action $u$ is given by the expectation of the population gain (15) with respect to $\mathfrak{Y}_t$ , i.e.

$$G_u = \mathbb{E}_{\mathfrak{Y}_t}\left[G_u^{\mathfrak{Y}_t}\right] \quad (16a)$$
$$= \sum_{h\in H_t}\sum_{n\geq 0} c_t(h) c^{\mathrm{a}}_t(n) G^{h,n}_u. \quad (16b)$$

Substituting (14) and (15) in (16a) yields the alternative form

$$G_u = \mathbb{E}_{\mathfrak{Y}_t}\Bigg[\sum_{Z\subseteq\mathbf{Z}_u}\Bigg[\sum_{\boldsymbol{h}\in\mathrm{Adm}_Z(\mathfrak{Y}_t)} \hat{P}^{\boldsymbol{a}}_u G^{\boldsymbol{a}}_u\Bigg]\Bigg], \quad (17)$$

where $\boldsymbol{a}$ consistently refer to $(h, n, \boldsymbol{h})$ where $(h, n)$ can be seen as a realisation of $\mathfrak{Y}_t$.

### C. Factorised expression of population gain $G_u$

We wish to factorise the expression (17) of the information gain $G_u$ with respect to the individual gains (9) in order to isolate the contribution of a specific track. In order to do so, we introduce the reduced population $\mathfrak{Y}^y_t$, where $y \in \bar{Y}_{t-1}$, as the estimated population $\mathfrak{Y}_t$ from which the target represented by $y$ is excluded. Realisations of $\mathfrak{Y}^y_t$ are of the form:

- $(h \setminus \{y\}, n)$, if $y \in Y_{t-1}$,
- $(h, n-1)$, if $y = \boldsymbol{\phi}_{t-1}$,

where $(h, n)$ is a realisation of $\mathfrak{Y}_t$. We can then factorise the population gain (16) as follows:

*Proposition 1:* The expected information gain $G_u$ for any action $u \in U_t$ is given by the sum

$$G_u = G^{\mathrm{d}}_u + G^{\mathrm{a}}_u, \quad (18)$$

where

$$G^{\mathrm{d}}_u = \sum_{z\in\bar{\mathbf{Z}}_u}\sum_{y\in Y_{t-1}} p^y_t(g_u(z,\cdot)) Q^{y,z}_u G^{y:z}_u, \quad (19)$$
$$G^{\mathrm{a}}_u = \sum_{z\in\mathbf{Z}_u} p^{\mathrm{a}}_t(g_u(z,\cdot)) Q^{\boldsymbol{\phi}_{t-1},z}_u G^{\mathrm{a}:z}_u, \quad (20)$$

with

$$Q^{y,z}_u = \mathbb{E}_{\mathfrak{Y}^y_t}\Bigg[\sum_{Z\subseteq\mathbf{Z}_u}\Bigg[\sum_{\boldsymbol{h}\in\mathrm{Adm}_{Z\setminus\{z\}}(\mathfrak{Y}^y_t)} \hat{P}^{\boldsymbol{a}}_u\Bigg]\Bigg]. \quad (21)$$

That is, the information gain $G_u$ is a linear combination of the individual gains for the updated tracks $\hat{Y}_u$, whether they characterise a previously detected target (19) or an appearing one (20).

We see in (19) that the contribution of each individual gain is weighted by *a)* $p^y_t(g_u(z,\cdot))$, i.e. the probability of the single-target/single-observation association leading to the updated track, and *b)* $Q^{y,z}_u$, i.e. the probability of the data association between the remaining population and observations. The same remark applies to the appearing targets in (20).

### D. Region-specific and/or track-specific information gain

The analysis carried in the previous section can be easily made specific to a given region of the state space *and/or* to a given set of targets For each track $y \in \bar{Y}_{t-1}$, we consider a measurable mapping $f^y : \bar{\mathbf{X}}_t \to \bar{\mathbf{X}}_t$, to be determined later, and we consider the probability measures in $\mathcal{P}(\bar{\mathbf{X}}_t)$

$$\begin{cases} f^y_*(p^y_t), f^y_*(p^{y:z}_u), & y\in Y_{t-1}, z\in\bar{\mathbf{Z}}_u, \\ f^{\boldsymbol{\phi}_{t-1}}_*(p^{\mathrm{a}}_t), f^{\boldsymbol{\phi}_{t-1}}_*(p^{\mathrm{a}:z}_u), & z\in\mathbf{Z}_u, \end{cases} \quad (22)$$

where $f_*(p)$ denotes the image of the measure $p$ under $f$ (see Section 3.6. in [14]), or pushforward measure, defined by

$$(f_*(p))(\cdot) = p(f^{-1}(\cdot)).$$

A careful choice of the mappings $f^y$ then allows an emphasis of the control on a desired region of the state space and/or a desired set of targets when the pushforward probabilities are substituted to the original ones in the construction of the individual gains (9).

*1) Example 1: region-specific information gain:* Assume that $B \in \mathbf{B}_{\mathbf{X}}$ denotes a region of special interest for the current sensor management decision – e.g. the immediate surroundings of a strategic building which is the focus of the surveillance activity. Consider a common mapping $f$ for all the tracks $y \in \bar{Y}_{t-1}$ and such that:

$$f(x) = \begin{cases} x, & \text{if } x\in B, \\ \psi, & \text{if } x\notin B. \end{cases} \quad (23)$$

Now, consider an updated track of the form "$y : z$" (see Section III-A1). Then the individual gain (9), when applied to the pushforward probabilities (22), becomes

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \Big[ \int_B \big[p_t^y(x)\big]^\alpha \big[p_u^{y:z}(x)\big]^{1-\alpha} \mu(\mathrm{d}x) + \big[p_t^y(\bar{\mathbf{X}}_t \setminus B)\big]^\alpha \big[p_u^{y:z}(\bar{\mathbf{X}}_t \setminus B)\big]^{1-\alpha} \Big]. \quad (24)$$

We see that the information gain is now focussed on the point-wise variations from $p_t^y$ to $p_u^{y:z}$ inside the region of interest $B$ *only*. Moreover, the gain is zero if and only if the (possibly empty) observation $z$ carried no additional information on the target regarding:

- its localization in $B$, since $p_t^y = p_u^{y:z}$ on $B$, and
- its presence in $B$, since $p_t^y(\bar{\mathbf{X}}_t \setminus B) = p_u^{y:z}(\bar{\mathbf{X}}_t \setminus B)$.

A similar reasoning applies for new tracks of the form "$\mathrm{a} : z$" (see Section III-A2).

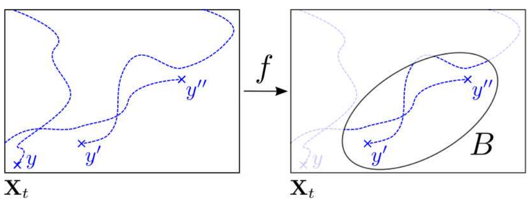


Fig. 3. Information gain assessed within $B$ only.

*2) Example 2: track-specific information gain:* Assume that $Y \subseteq Y_{t-1}$ denotes a set of tracks of special interest driving the current sensor decision – e.g. representing targets with a significant level of threat. Consider the mappings $f^y$ such that:

$$f^y(x) = \begin{cases} x, & \text{if} \quad y \in Y, \\ \psi, & \text{if} \quad y \notin Y. \end{cases} \quad (25)$$

Now, consider an updated track of the form "$y : z$" (see Section III-A1). If $y \in Y$ then the individual gain (9) remains unchanged when applied to the pushforward probabilities (22). On the other hand, if $y \notin Y$ it becomes

$$G_u^{y:z} = \frac{1}{\alpha - 1} \log \Big[ \big[p_t^y(\bar{\mathbf{X}}_t)\big]^\alpha \big[p_u^{y:z}(\bar{\mathbf{X}}_t)\big]^{1-\alpha} \Big] = 0.$$

In other words, only the information gain on the targets with specific interest is involved in the decision policy. A similar reasoning applies for new tracks of the form "$\mathrm{a} : z$" (see Section III-A2).

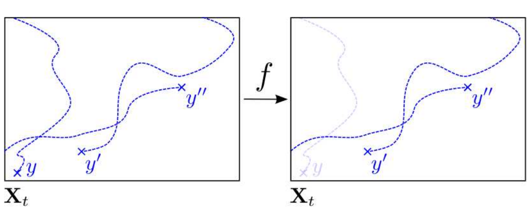


Fig. 4. Information gain assessed for tracks $y'$, $y''$ only.

## IV. Conclusion

This paper addresses the problem of closed-loop sensor management for multi-object filtering solutions using stochastic populations of independent targets. The derivation of a reward function for each possible sensor action is provided in a principled and general way, and further expressed as a linear combination of the expected information gains, determined with the Rényi divergence, for each track composing the estimated target population. The proposed reward function can be focussed on specific regions of the scene and/or specific targets, and this is illustrated on two simple examples. Filtering approximations will be subsequently explored for an efficient computation of the weights involved in the linear combination.

## Acknowledgements

This work was supported by the Engineering and Physical Sciences Research Council (EPSRC) Platform Grant (EP/J015180/1), the MOD University Defence Research Centre on Signal Processing (UDRC) Phase 1 (EP/H011900/1) and Phase 2 (EP/K014227/1). Jeremie Houssineau has a PhD scholarship sponsored by DCNS and a tuition fee scholarship sponsored by Heriot-Watt University.

## References

[1] B.-N. Vo, B.-T. Vo, and D. E. Clark. Bayesian Multiple Target Filtering Using Random Finite Sets. In M. Mallick, V. Krishnamurthy, and B-N. Vo, editors, *Integrated Tracking, Classification, and Sensor Management*, chapter 3, pages 75–125. Wiley, 2013.

[2] D. A. Castanón and L. Carin. Stochastic control theory for sensor management. In *Foundations and applications of sensor management*, pages 7–32. Springer, 2008.

[3] A. O. Hero III, C. M. Kreucher, and D. Blatt. Information theoretic approaches to sensor management. In *Foundations and applications of sensor management*, pages 33–57. Springer, 2008.

[4] R. P. S. Mahler. *Statistical Multisource-Multitarget Information Fusion*. Artech House, 2007.

[5] R. P. S. Mahler. Objective Functions for Bayesian Control-Theoretic Sensor Management, I: Multitarget First-Moment Approximation. In *Aerospace Conference, IEEE Proceedings on*, volume 4, pages 1905–1923, March 2003.

[6] B. Ristic, B.-N. Vo, and D. E. Clark. A note on the Reward Function for PHD Filters with Sensor Control. *Aerospace and Electronic Systems, IEEE Transactions on*, 47(2):1521–1529, April 2011.

[7] A. K. Gostar, R. Hoseinnezhad, A. Bab-Hadiashar, and B.-T. Vo. Control of sensor with unknown clutter and detection profile using Multi-Bernoulli filter. In *Information Fusion, Proceedings of the 16th International Conference on*, pages 1021–1028. IEEE, 2013.

[8] A. K. Gostar, R. Hoseinnezhad, and Alireza Bab-Hadiashar. Multi-Bernoulli sensor control for multi-target tracking. In *Intelligent Sensors, Sensor Networks and Information Processing, 2013 IEEE Eighth International Conference on*, pages 312–317. IEEE, 2013.

[9] H. G. Hoang. Control of a mobile sensor for multi-target tracking using multi-target/object Multi-Bernoulli filter. In *Control, Automation and Information Sciences (ICCAIS), 2012 International Conference on*, pages 7–12. IEEE, 2012.

[10] J. Houssineau and D. E. Clark. Representation and estimation of stochastic populations. Technical report, Heriot Watt University, 2014.

[11] A. O. Hero III. Sensor Management: Past, Present, and Future. *IEEE Sensor Journals*, 11(12):3064–3075, December 2011.

[12] A. Rényi. On measures of entropy and information. In *Fourth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 1961.

[13] P. Del Moral and A. Doucet. Particle methods: An introduction with applications. Technical report, INRIA, 2009.

[14] V. I. Bogachev. *Measure Theory*, volume 1. Springer, 2007.