

A theoretical look at information-driven sensor management criteria

Edson Hiroshi Aoki¹, Arunabha Bagchi¹, Pranab Mandal¹, and Yvo Boers²

¹Department of Applied Mathematics, University of Twente, Enschede, The Netherlands.

²Thales Nederland B. V., Hengelo, The Netherlands.

¹[1{e.h.aoki, a.bagchi, p.k.mandal}@ewi.utwente.nl](mailto:{e.h.aoki, a.bagchi, p.k.mandal}@ewi.utwente.nl)

²[2yvo.boers@nl.thalesgroup.com](mailto:yvo.boers@nl.thalesgroup.com)

Abstract—In sensor management, the usefulness of information theoretic measures seems to be validated by a large number of empirical studies, but theoretical justification presented until so far, both for selection of the measure and for the use of information-driven sensor management itself, still seems inconclusive, conflicting, or debatable.

In this paper, we suggest that information-driven sensor management may be justified on the basis of uncertainty reduction rather than information gain.

We subsequently identify that, due to well-known relationships between Shannon entropy, mutual information and Kullback-Leibler (KL) divergence, for sensor management purposes using the Kullback-Leibler (KL) divergence (a measure of information gain; thus a relative measure) is exactly the same as using the Shannon entropy (a measure of uncertainty; an absolute measure).

This is also used to demonstrate that, if uncertainty reduction is desirable, the asymmetry of the KL divergence is not relevant to the sensor management problem. Finally, we show some counterpoints to some arguments for replacing the KL divergence with the more general Rényi divergences.

Keywords: Sensor management, information theory, entropy, Kullback-Leibler divergence, Rényi divergence.

I. INTRODUCTION

Sensor management is a control problem associated with partially observed systems, where the control action aims to influence the generation of observations (direct feedthrough) and not the state of the system, generally with the goal of obtaining the best possible estimation quality of the state given limited sensing resources.

Typically the sensor management problem is formulated in terms of minimization of a risk function related to the error between the true state and the estimated state; this is the so-called “task-driven” sensor management. An alternative is instead attempting to improve (in some sense) the “information content” of the distribution. This “information-driven” sensor management consists of choosing the control decision that maximizes some notion of information gain (or, alternatively, minimizes some notion of uncertainty).

Hintz and McVey [1], [2] were the first to suggest using information theory in a problem related to sensor management and state estimation, followed by Manyika and Durrant-Whyte [3], who considered expected information gain in sensor management and data fusion problems. The idea of using the

Kullback-Leibler (KL) divergence for sensor management appeared in the works of Schmaedeke and Kastella [4], Kastella [5], and Mahler [6]. Doucet, Andrieu and Thomas [7] provided a particle filter implementation of sensor management based on KL divergence, which could be used for general non-linear systems.

The idea of using the more general Rényi divergence (or α -divergence), instead of the KL divergence, was introduced by Kreucher, Kastella and Hero [8] and demonstrated for a multitarget tracking problem. In this initial work, the authors do not provide a particular reason for using the α -divergence instead of the KL divergence, except that it gives an extra freedom of choosing the parameter α , which could, in principle, emphasize certain parts of the distribution functions.

In [9], Kreucher, Kastella and Hero suggest using either $\alpha = 0.5$ (which corresponds to the Hellinger affinity) or $\alpha = 1$ (which corresponds to the KL divergence), based on another work of them [10], where an empirical study on resolution of clusters for a georegistration problem was done. The authors also suggest the use of the value $\alpha = 0.5$ when the prior and posterior densities are similar. This is based on an asymptotic analysis of the Chernoff exponent [11], used in hypothesis testing between two probability densities.

However, Aughenbaugh and La Cour [12] observed that the supposed superior discrimination capability of the α -divergence with $\alpha = 0.5$ does not necessarily correspond to our intuitive interpretation of information gain. Through analysis of a few examples, they observed that using $\alpha = 0.5$ for sensor management, as opposed to using the KL divergence, resulted in actions that emphasized morphological changes on the distribution (such as rotation, or translation of modes for multi-modal distributions) which did not necessarily correspond to desirable sensor management properties.

Two subsequent works of Kreucher, Kastella and Hero [13], [14] made a strong theoretical argument in favor of information-driven sensor management, by claiming that the expected value of arbitrary risk functions is sandwiched between functions of two marginalized Rényi divergences; this would make a criterion based on the α -divergences a “near-universal” proxy for task-driven sensor management. This argument was rebutted in Aoki et al. [15], which shows that on basis of these sandwich bounds one cannot affect task-driven

goal functions by maximizing or minimizing α -divergences.

A very recent work of Aughenbaugh and La Cour [16] also compares information-theoretic measures for sensor management applied to the problem of tracking a maneuvering target observed by multistatic sensors. In this work they state the well-known relationship [17], [18] between Shannon entropy, mutual information and expected KL divergence, but the comparison between criteria is done empirically. Finally, Ristic, Vo and Clark [19] describe the implementation of α -divergences for PHD filters, and also present an empirical comparison between these and other criteria.

This work is organized as follows. Section II briefly describes the sensor management problem, and also describes the task-driven and information-driven approaches. Section III analyses entropy and possible reasons for its application in sensor management. Section IV discusses the theoretical justification of the Kullback-Leibler and Rényi divergences as sensor management criteria. Section V draws the conclusions of this work.

II. THE SENSOR MANAGEMENT PROBLEM

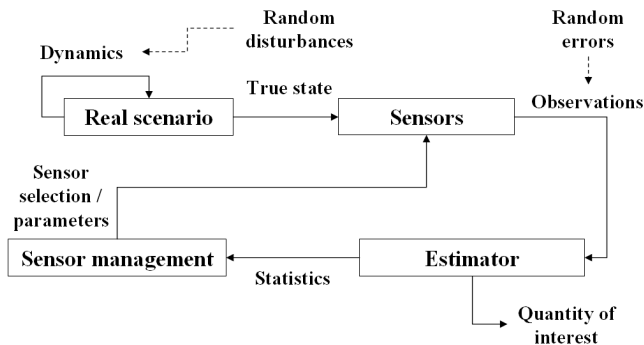


Figure 1. Sensor management as a stochastic control problem

A representation of sensor management as a stochastic control problem is shown on Figure 1. We consider a (static or dynamic) scenario described by a state X , observed by a measurement device composed of one or more sensors, with the sensor observations Y corrupted by random errors. These measurements are used as input to an estimator, which obtains an estimate $\hat{\Psi}$ of a quantity of interest Ψ , in general a function of X .

A feedback occurs through a sensor management device, which uses $\hat{\Psi}$ or other statistics computed by the estimator to select a control decision (or “sensing action”) U that affects the generation of subsequent observations. Typically, U is chosen by minimization (or maximization) of the expectation of a risk (or reward) function $\gamma(X, Y, U)$.

As we can see from Figure 1, the difference between sensor management and the standard control problem is that, in the first, the control decision aims to affect the generation of observations (i.e. direct feedthrough input), while in the second, it aims to affect the true state.

More generally, we can also perform “long-term” sensor management, i.e. to make our control decision considering

what will happen in a time horizon $k_0 + 1, \dots, k_0 + H$. In this case, the goal function γ depends on X_{k_0+1, \dots, k_0+H} , Y_{k_0+1, \dots, k_0+H} , and U_{k_0+1, \dots, k_0+H} . Each decision U_k is not modelled as deterministic but as a random variable which may depend on information $Y_{k_0+1, \dots, k-1}$ and $U_{k_0+1, \dots, k-1}$. Therefore, our goal is not actually to determine a realization u_k for U_k but instead a control law η_k , where $U_k = \eta_k(Y_{k_0+1, \dots, k-1}, U_{k_0+1, \dots, k-1})$. For the sake of simplicity we will not consider this more general formulation, although everything discussed in this paper also applies to it.

From now on, we are going to use the notation $\mu_{A|B}$ to refer to the distribution of a random variable A conditioned on a random variable or control decision B . In particular, we will refer to μ_X and $\mu_{X|Y,U}$ as prior and posterior distributions of X respectively. For convenience, the time index and the conditioning on past measurements and control decisions will be implicit on all distributions and densities.

A. Task-driven sensor management

Let $\epsilon(\Psi, \hat{\Psi})$ be a performance metric for our estimator, corresponding to some measure of error between the estimated quantity $\hat{\Psi}$ and the true quantity Ψ . For instance, the square error of a scalar quantity would be given by $\epsilon(\Psi, \hat{\Psi}) = (\Psi - \hat{\Psi})^2$.

In task-driven sensor management, we directly attempt to optimize our performance metric, i.e. the control decision is chosen by the optimization

$$\arg \min_U E_{\mu_{X,Y|U}} [\epsilon(\Psi, \hat{\Psi})]. \quad (1)$$

Naturally, we can define task-driven criteria that do not precisely have form (1), such as when we use a performance metric that is not function of the true state (for instance, the maximum posterior probability suggested in [13]). While performing task-driven sensor management is intuitive and seemingly straightforward, it can be impractical in some situations.

First, as also remarked by Kreucher, Kastella and Hero [13], it may be difficult to define a metric of the form $\epsilon(\Psi, \hat{\Psi})$ that meaningfully measures performance of our estimator, for instance when we have multiple goals with subjective importance. This also happens when our goals are defined in different dimensions or state spaces.

Second, if obtaining our performance metric already incurs a significant computation cost, performing sensor management based on it may be computationally unfeasible. This is because performing sensor management is much more expensive than computing performance metrics, since according to (1), we also need to integrate over X and Y , and the goal function must be evaluated for every possible control decision u . Also, although we usually can compute a performance metric off-line, evaluating goal functions for sensor management must be made on-line.

B. Information-driven sensor management

One alternative to task-driven sensor management is to define a goal function based on the distribution of X itself, not on the estimate or the true state. In information-driven

sensor management, we attempt to maximize the “information content” of the posterior, i.e. its capacity of yielding (in some sense) useful information to the operator, rather than attempting to maximize the quality of a particular estimate.

For instance, we may obtain the control decision by

$$\arg \min_U E_{\mu_{Y|U}} [f(\mu_{X|Y,U})] \quad (2)$$

where f is some measure of information content of the posterior distribution. One common choice of f is the Shannon entropy (or one of its generalisations, including the Rényi entropy). Alternatively, instead of just looking at the posterior distribution, we may attempt to maximize some notion of information gain between prior and posterior distributions. In this case, we have

$$\arg \max_U E_{\mu_{Y|U}} [f(\mu_{X|Y,U}, \mu_X)]. \quad (3)$$

Examples of f include the Kullback-Leibler divergence and its generalisations, including the Rényi divergence. In a way, we can say that task-driven sensor management conforms to the paradigm that the purpose of a filter is to yield point estimates of the quantities of interest. Information-driven sensor management, on the other hand, conforms to the paradigm that the purpose of a filter is to compute accurate statistical descriptions (such as posterior densities) of these quantities.

III. THE MEANING OF ENTROPY IN SENSOR MANAGEMENT

Most of the previous studies on information-theoretic sensor management have focused on the information gain (in terms of KL or α -divergences) between prior and posterior distribution. However, although commendable efforts [13], [14] were made, these works have not been able to find a clear, theoretically sound relationship between information gain and performance metrics used in estimation.

Instead of information gain, a few works ([18], [20]) have hinted at looking at the amount of “overall uncertainty” of the posterior distribution, i.e., its absolute entropy. We therefore find worthwhile to take an extra look at entropy and its usefulness in sensor management.

The Shannon entropy of a distribution μ_X of some random variable X , with respect to a dominating σ -finite reference measure ρ , is defined by

$$H(X) \triangleq - \int_{\mathcal{X}} p_X(x) \log p_X(x) \rho(dx) \quad (4)$$

where $p_X = \frac{d\mu_X}{d\rho}$ is the density, i.e. the Radon-Nikodym derivative, of μ_X with respect to ρ and \mathcal{X} is the support of p_X .

For a sequence of i.i.d. discrete random variables, the Shannon source coding theorem relates entropy to the minimum necessary number of bits for lossless data compression of the sequence. In other words, the Shannon entropy corresponds to a notion of the amount of “uncertainty” contained in a distribution, since the less predictable the symbols of the sequence are, the greater the risk is of data loss by not completely encoding every symbol. Some relatively similar results can be obtained for continuous random variables (see [17]).

A. Entropy for Gaussian and Gaussian mixture distributions

The Shannon source coding theorem does not show, however, a clear link between entropy and estimation performance. Therefore, to simplify our analysis, we will, for the time being, consider only Gaussian and Gaussian mixture distributions.

Let us consider a continuous random variable X with density p_X . If $p_X(x) = \mathcal{N}(x; \hat{x}, P)$, we have the following well-known relationship between the Shannon entropy and the covariance of the MMSE estimate

$$H(X) = \frac{n_x}{2} \log 2\pi e + \frac{1}{2} \log |P|. \quad (5)$$

where n_x is the dimension of the state space and P is the covariance. Therefore in the Gaussian case the entropy is monotonically increasing w.r.t. the covariance determinant.

Now, for $p_X(x) = \sum_{i=1}^m w_i \mathcal{N}(x; \hat{x}_i; P_i)$, the following relationship [21] holds

$$\begin{aligned} & - \sum_{i=1}^m w_i \log \left(\sum_{j=1}^m w_j \mathcal{N}(\hat{x}_i; \hat{x}_j, P_i + P_j) \right) \\ & \leq H(X) \leq \\ & \sum_{i=1}^m w_i \left(-\log w_i + \frac{n_x}{2} \log 2\pi e + \frac{1}{2} \log |P_i| \right). \end{aligned} \quad (6)$$

We recall that the covariance of a Gaussian mixture with mean \hat{x} is given by

$$P = \sum_{i=1}^m w_i (P_i + (\hat{x}_i - \hat{x})(\hat{x}_i - \hat{x})'). \quad (7)$$

As we can see from (7), if we arbitrarily increase the distance between the means of the components of the mixture, we may also arbitrarily increase P . On the other hand, by looking at (6), we see that the lower bound of $H(X)$ depends on both means and covariances of individual components of the mixture, but the upper bound depends only on covariances. So, if we arbitrarily increase the distance between modes, we can at most make $H(X)$ closer to the upper bound.

Therefore, if for a Gaussian distribution we assume that covariance corresponds to our intuitive notion of “uncertainty”, for a Gaussian mixture with sufficient separation between the modes, entropy corresponds to the amount of uncertainty around each component, without regard of the relative position of the components.

Can we extend such relations to information-driven sensor management criteria of form (2)? Let $p_{X|Y,U}(x|y, u) = \mathcal{N}(x; \hat{x}(y, u), P(y, u))$ be the density of the posterior $\mu_{X|Y,U}$, and let $H(X|y, u)$ be the entropy of $\mu_{X|Y,U}$ (for fixed y and u). By applying Jensen’s inequality to (5), we obtain

$$H_c(X|Y, u) \leq \frac{n_x}{2} \log 2\pi e + \frac{1}{2} \log E_{\mu_{Y|U}} [|P(Y, u)|] \quad (8)$$

where $H_c(X|Y, u) \triangleq E_{\mu_{Y|U}} [H(X|Y, u)]$ is the so-called conditional entropy of $\mu_{X|Y,U}$. Similarly, let $p_{X|Y,U}(x|y, u) =$

$\sum_{i=1}^m w_i \mathcal{N}(x; \hat{x}_i(y, u), P_i(y, u))$. Then, from (6), we have

$$H_c(X|Y, u) \leq \sum_{i=1}^m w_i \left(-\log w_i + \frac{n_x}{2} \log 2\pi e + \frac{1}{2} \log E_{\mu_{Y|U}}[|P_i(Y, u)|] \right). \quad (9)$$

Result (8) shows that, in the Gaussian case, by minimizing the conditional entropy one also minimizes a lower bound on the expectation of the determinant of the MMSE covariance. Similarly, result (9) shows that in the Gaussian mixture case, minimization of the conditional entropy causes a similar effect on the covariances of individual mixture components.

B. Entropy as a sensor management criterion

We have shown that the minimum entropy and minimum covariance sensor management criteria have some sort of equivalence in the Gaussian posterior case. But in the Gaussian mixture posterior case, it is not clear why we would consider using entropy. For instance, let us assume that $X = [X^{(1)} \dots X^{(n_x)}]'$ and the corresponding estimate (MMSE or otherwise) is given by $\hat{x} = [\hat{x}^{(1)} \dots \hat{x}^{(n_x)}]'$. Let us then consider that we have chosen the sum of the square errors of \hat{x} as performance metric. Then according to (1), the corresponding task-driven criterion is given by

$$\begin{aligned} E_{\mu_{X,Y|U}} \left[\sum_{j=1}^{n_x} \left(X^{(j)} - \hat{x}^{(j)}(Y, u) \right)^2 \right] \\ = E_{\mu_{Y|U}} \left[\sum_{j=1}^{n_x} E_{\mu_{X|Y,U}} \left[\left(X^{(j)} - \hat{x}^{(j)}(Y, u) \right)^2 \right] \right] \\ = E_{\mu_{Y|U}} [\text{tr}(P(Y, u))] \end{aligned} \quad (10)$$

where P is the covariance of \hat{x} . Since (10) is the criterion that directly optimizes our performance metric, one may ask why we should use anything else than it, such as entropy.

The flaw of this reasoning is assuming that we have a representative estimate in the first place, and also assuming that we have a corresponding representative performance metric. This is generally true for mono-modal distributions, but for multi-modal distributions (including Gaussian mixtures with sufficient component separation), finding a suitable estimate-metric pair can be very tricky. An example is shown on Figure 2.

Distribution B has smaller variance, but higher entropy than distribution A. Therefore, if we use minimization of expected squared errors (i.e. variance) as sensor management criterion, our scheme would prefer an action that leads to the distribution B, rather than distribution A. However, which distribution gives us more information about the actual state? By looking at the distribution A, we know that the state is likely to be very close to point $x = -5$, and if that is not the case, to point $x = 5$. In contrast, by looking at the distribution B we just have a vague idea of where the real state is!




Figure 2. Two Gaussian mixture PDFs, the first with variance 16.25 (with respect to the MMSE estimate) and entropy within the interval [1.07,1.23], and the second with variance 7.84 and entropy within [2.32,2.79]

The reason is that the squared errors of the MMSE estimate are not a good performance metric in this case. Due to multimodality, the MMSE estimate is not representative as it just gives a mean of the modes, without giving much hint on where the true state can actually be. By looking at (7), we can see that the expected sum of square errors (10) includes a measure of separation of the estimate from the actual modes, or in other words, a measure of separation between the modes. The MAP estimate would be probably more representative in this case, but its expected square errors will also strongly depend on the separation between the modes.

On the other hand, by looking at (6), we can see that entropy depends mostly on the covariance of each component, rather than the separation between the components. So, when the entropy becomes low enough, we will typically be able to accurately obtain a list of possible localizations of the state (each corresponding to one of the components of the mixture), although we will often not be able to determine which is the correct component.

Although we considered only the Gaussian and Gaussian mixture cases, we may expect that what we discussed here is valid, obviously in different degrees, for other kinds of distributions. By adopting the same terms used by Aughenbaugh and La Cour [12], we could say that entropy is more closely related to the “overall uncertainty” of a multi-modal distribution, where covariance is more closely related to its amount of “ambiguity”.

As we can see, the so-called task driven sensor management relies on the idea that our filter must produce a point estimate as output, and that we can meaningfully measure the performance of this point estimate if ground truth information is available. However, if such conditions are not met, we may decide to look directly at the posterior distribution (i.e. the statistical information about the quantities that we want to estimate), and in this case a criterion based on entropy would

make more sense.

C. An example: tracking of closely spaced targets

An example of a multi-modal problem is joint tracking of closely spaced targets, where, according to Boers, Sviestins and Driessen [22], we typically have a mode for the correct labeling of targets and other modes corresponding to incorrect labelings. This “mixed labelling” situation arises when two or more targets move close to each other for some time, and thereafter they separate.

Boers et al. [20] have analyzed this phenomenon in detail for the case of two targets. They observed that when two confused targets separate, entropy starts decreasing (as the individual target estimates become more accurate, since there is no more confusion on the origin of measurements), but the determinant of the covariance of the MMSE estimate starts increasing (as the two modes, i.e. the two alternate labelings, go further from each other). This is an empirical example of how the behavior that we identified for Gaussian mixtures can be seen on general multi-modal problems.

The apparent consequence is that using the corresponding square errors as criterion (i.e. using (10)) will result in a risk function that may be dominated by the target separation (i.e. the distance between the modes), rather than by the individual target states’ errors (i.e. the uncertainty around each mode).

In this case, using an information-driven criterion based on entropy would make sense, since as we have seen, entropy seems to be associated with the amount of “overall uncertainty”, regardless of the presence of multi-modality. An alternative would be using a task-driven criterion based on the OSPA metric defined in [23], but we note that this approach may easily have prohibitive computational cost.

D. The Rényi entropy

By taking the same Fadeev’s postulates used to characterize the Shannon entropy, one can define a generalisation of the Shannon entropy [24], which correspond to the Rényi entropies

$$H_\alpha(X) \triangleq \frac{1}{1-\alpha} \log \int_{\mathcal{X}} p_X^\alpha(x) \rho(dx) \quad (11)$$

where the parameter α may be used to give more or less emphasis to low probability regions (“tails”) of the distribution. The Rényi entropy and the MMSE covariance have a relationship similar to (5), so we assume that most of what we discussed for Shannon entropies is also valid for general Rényi entropies. Finding situations where it is advantageous to use general Rényi entropies instead of the Shannon entropy in sensor management is a topic for future discussion.

IV. ON THEORETICAL JUSTIFICATION OF KULLBACK-LEIBLER AND RÉNYI DIVERGENCES

A. Some theoretical observations about KL and α -divergences

The relative Shannon entropy, or Kullback-Leibler (KL) divergence is a measure of difference between two distributions. Consider a pair of distributions μ and ν which respectively

admit densities p and q with respect to a dominating σ -finite measure ρ . The KL divergence from μ to ν (or alternatively, from p to q) is given by¹

$$D_{\text{KL}}(p||q) \triangleq \int p(x) \log \frac{p(x)}{q(x)} \rho(dx). \quad (12)$$

where we apply the conventions $\log \frac{p(x)}{q(x)} = 0$ for $p(x) = 0$ and $q(x) = 0$, and $a/0 = \infty$ for $a > 0$. Note that the measure is asymmetric in the sense that $D_{\text{KL}}(p||q) \neq D_{\text{KL}}(q||p)$.

Similarly, the Rényi divergence or α -divergence is a generalisation of the KL divergence in the sense that it satisfies a set of postulates that characterize the KL divergence. The α -divergence from μ to ν (or alternatively, from p to q) is given by

$$D_\alpha(p||q) \triangleq \frac{1}{\alpha-1} \log \int p^\alpha(x) q^{1-\alpha}(x) \rho(dx). \quad (13)$$

where we apply the conventions $p^\alpha(x) q^{1-\alpha}(x) = 0$ for $p(x) = 0$ and $q(x) = 0$, and $a/0 = \infty$ for $a > 0$. D_0 and D_1 are defined using the limits from right and left respectively, which makes D_1 the same as the D_{KL} . $D_{0.5}$ has the special property that it is a true metric, in the sense that it is symmetric and obeys the triangle inequality.

In information-driven sensor management, these divergences are used on criteria of the form (3), i.e. they are considered to represent a notion of information gain between the prior μ_X and the posterior $\mu_{X|Y,U}$. Thus, we would be interested in choosing the sensing action that maximizes these divergences. Since KL divergences and α -divergences are asymmetric (except for $\alpha = 0.5$), one may ask which order of the arguments shall be used, i.e. whether we should maximize

$$\gamma_1 \triangleq E_{\mu_{Y|U}} [D(p_{X|Y,U} || p_X)] \quad (14)$$

or

$$\gamma_2 \triangleq E_{\mu_{Y|U}} [D(p_X || p_{X|Y,U})] \quad (15)$$

where D denotes an unspecified divergence measure. For instance, in [7], it is suggested that we use the average of γ_1 and γ_2 . In [25], this asymmetry was stated as a reason to use the $\alpha = 0.5$ divergence instead of the KL divergence, as a criterion based on the KL divergence would not be globally consistent (i.e. we may get different decisions if we exchange the order of the arguments).

A counterpoint to this argument lies in the well-known relation between the expected KL divergence and conditional entropy. If we have $p_X = p_{X|U}$, which holds under a few assumptions, it is easy to see that only γ_1 corresponds to the mutual information $I(X; Y)$ between state X and observation

¹The KL divergence can also be defined for distributions which do not admit a density w.r.t. a common reference measure, but such highly general definition is not relevant for this work.

Y. In this case, we have

$$\begin{aligned}
I(X; Y) &\triangleq \int p_{X,Y|U}(x, y|u) \log \frac{p_{X,Y|U}(x, y|u)}{p_{X|U}(x|u)p_{Y|U}(y|u)} \\
&\quad \times \rho \otimes \rho(dx \times dy) \\
&= \int p_{Y|U}(y|u) \int p_{X|Y,U}(x|y, u) \log \frac{p_{X|Y,U}(x|y, u)}{p_X(x)} \\
&\quad \times \rho(dx) \rho(dy) \\
&= E_{\mu_{Y|U}}[D_{\text{KL}}(p_{X|Y,U} \| p_X)]. \tag{16}
\end{aligned}$$

Moreover, $I(X; Y)$ can be rewritten as

$$\begin{aligned}
I(X; Y) &= \int p_{Y|U}(y|u) \int p_{X|Y,U}(x|y, u) \log p_{X|Y,U}(x|y, u) \\
&\quad \times \rho(dx) \rho(dy) - \int p_X(x) \log p_X(x) \rho(dx) \\
&= -E_{\mu_{Y|U}}[H(X|Y, u)] + H(X) \\
&= -H_c(X|Y, u) + H(X). \tag{17}
\end{aligned}$$

Since $H(X)$ does not depend on u , relations (16) and (17) imply that maximization of γ_1 as criterion is equivalent to minimization of conditional entropy. This means that *for sensor management purposes minimizing conditional entropy and maximizing expected KL divergence (or equivalently, mutual information) will lead to exactly identical sensing actions*. To the best of our knowledge, this equivalence has not been explicitly stated elsewhere, although relations (16) and (17) have been known for quite a while, see e.g. [16]–[18]. Note that the equivalence is ignored in some works such as [26], where the two criteria are empirically compared (and even different results are obtained, due to the use of suboptimal heuristics). Note also that this equivalence assumes that $p_X = p_{X|U}$, which does not hold in certain cases (such as the problem of tracking smart targets, where the targets are assumed to be aware of the sensing actions and may react to them).

In contrast, γ_2 does not have, to the best of our knowledge, this equivalence property. Hence, if, following our discussion in Section III, we consider that minimization of entropy is desirable, the asymmetry of the KL divergence with respect to its arguments is irrelevant because there is a “correct” order of arguments to be used!

Can we get a similar relationship for the Rényi divergence? Let $H_\alpha(X)$ and $H_\alpha(X|y, u)$ be the Rényi entropies of μ_X and $\mu_{X|Y,U}$ respectively. Let $I_\alpha(X; Y)$ be defined as $I_\alpha(X; Y) \triangleq H_\alpha(X) - E_{\mu_{Y|U}}[H_\alpha(X|Y, u)]$. Observe that

$$\begin{aligned}
I_\alpha(X; Y) &= H_\alpha(X) - E_{\mu_{Y|U}}[H_\alpha(X|Y, u)] \\
&= -\frac{1}{\alpha-1} \log \int p_X^\alpha(x) \rho(dx) + \frac{1}{\alpha-1} \int p_{Y|U}(y|u) \\
&\quad \times \log \left(\int p_{X|Y,U}^\alpha(x|y, u) \rho(dx) \right) \rho(dy). \tag{18}
\end{aligned}$$

On the other hand, the expected value of the Rényi divergence $D_\alpha(p_{X|Y,U} \| p_X)$ is given by

$$\begin{aligned}
E_{\mu_{Y|U}}[D_\alpha(p_{X|Y,U} \| p_X)] &= \frac{1}{\alpha-1} \int p_{Y|U}(y|u) \\
&\quad \times \log \left(\int p_{X|Y,U}^\alpha(x|y, u) p_X^{1-\alpha}(x) \rho(dx) \right) \rho(dy). \tag{19}
\end{aligned}$$

Although one can see similarities between expressions (18) and (19), they do not generally lead to the same results. For instance, let us choose two distributions μ_X and μ_Y of two discrete random variables X and Y respectively, with corresponding joint distribution μ_{XY} . Let these distributions be described by

$$\begin{aligned}
P(X=0) &= 0.3, & P(X=1) &= 0.7, \\
P(Y=0) &= 0.4, & P(Y=1) &= 0.6, \\
P(X=0, Y=0) &= 0.2, & P(X=0, Y=1) &= 0.1, \\
P(X=1, Y=0) &= 0.2, & P(X=1, Y=1) &= 0.5.
\end{aligned}$$

For these distributions, Figure 3 shows the difference $I_\alpha(X; Y) - E_{\mu_{Y|U}}[D_\alpha(p_{X|Y,U} \| p_X)]$ for varying values of α (using 2 as base for the logarithms).




Figure 3. Difference between $I_\alpha(X; Y)$ (which depends on the expected Rényi entropy) and the expected α -divergence for the chosen arbitrary distributions

As we may expect, the equality holds for $\alpha = 1$, in which case the Rényi entropy and the α -divergence become the Shannon entropy and the KL divergence respectively. Therefore, at least to the best of our knowledge, there is no equivalence between the maximum α -divergence and minimum Rényi entropy criteria. Some empirical results in literature [14], however, indicate that both criteria may behave similarly from a practical point of view.

B. A discussion on previous arguments for using the α -divergence instead of the KL divergence

It is easy to think that since the KL divergence is only a particular case of the Rényi divergence, one should use the more general form to enjoy the freedom of tuning the α parameter, and thus optimize performance in some sense.

We should keep in mind, however, that if a non-parameterized solution is optimal in some sense, adding an extra parameter may cause deviation from this optimal condition.

For instance, the KL divergence criterion, at least apparently, is the only Rényi divergence criterion which is equivalent to minimization of the corresponding entropy. Therefore, if we are interested on minimization of entropy, there is a strong reason for using the KL divergence instead of α -divergences for any value of α other than 1.

In this section, we discuss some arguments (beyond the asymmetry argument) used in supporting the use of the more general α -divergences instead of the KL divergence as criteria for sensor management:

- 1) *By using the Rényi divergence instead of the KL divergence, one can adjust the α parameter and give emphasis to specific parts of the distribution.*

This argument seems to be based on the fact that, by adjusting the α parameter of the Rényi entropy, one can compute a measure of uncertainty giving more emphasis to tails (i.e. regions with low probability) of the distribution.

This argument is debatable because as we have previously seen, there seems to be no obvious relationship between maximizing divergence and minimizing entropy, except for $\alpha = 1$. In fact, by looking at the Rényi divergence formula (13), it seems that if we attempt to give more emphasis to low probability regions of the posterior density (by decreasing α), we will automatically reduce the emphasis on low probability regions of the prior density, so interpreting this “emphasis” effect in an intuitive manner seems tricky to say the least.

- 2) *By using the Rényi divergence instead of the KL divergence, we may choose the value of α that maximizes the “discriminating capacity” between posterior and prior distributions; this would correspond to the “optimal” value of α to use.*

As remarked by Hero et al. [11], one can find a value of α for which the α -divergence gives the probability of error in hypothesis testing using the optimal Bayes test; in some sense, this would be the value of α that maximizes our capacity of discriminating the posterior and prior distributions. The problem is that there is no clear theoretical link between this “discriminating capacity” and sensor management performance. For instance, in [12], Aughenbaugh and La Cour verified that for Gaussian priors and posteriors, using the $\alpha = 0.5$ divergence (which should have more “discriminating capacity”) instead of the KL divergence lead to actions that emphasized rotations of the covariance matrix, and gave less emphasis to its reduction.

- 3) *Empirical results found in literature (e.g. [9], [27]) show that the Rényi divergence for some value of $\alpha \neq 1$ results in superior sensor management performance than the KL divergence.*

The problem of interpreting empirical results for non-conventional distributions should be evident from our

previous discussions. For instance, in [9], the position error of a state estimate is used as metric for a problem of multi-target tracking with frequent target crossings. But such problem is prone to generating multi-modal distributions; as we discussed in Section III-B, point estimates for multi-modal posteriors may be non-representative and also lead to non-representative performance metrics.

From that discussion, it seems reasonably difficult to justify replacing the KL divergence with the Rényi divergence with a simple argument of generalisation. We should also note that there are other possible generalisations of the KL divergence, including Arimoto α -divergencies, f -divergences and even the entropy difference described by Eq. (18). Interested readers may take a look at the excellent analysis of Liese and Vajda [28], which describes various types of divergences used in information theory, their main properties, and relationships.

V. CONCLUSIONS

In this paper we examined the information-driven approach to sensor management from a more theoretical point of view. We suggested that it is perhaps easier to find a justification for information-driven sensor management by looking at the entropy of the posterior, instead of looking at the information gain between the prior and the posterior. We also suggested that a criterion based on minimization of entropy could be particularly useful for multi-modal distributions, where we find difficulty on extracting point estimates and measuring their performance.

We also analyzed the Kullback-Leibler divergence and the Rényi divergence criteria commonly used in sensor management. We showed the equivalence between the KL divergence and the Shannon entropy criteria for sensor management purposes, which indicates that, even through the KL divergence is only a particular case of the Rényi divergence, the non-empirical justification presented until so far for replacing the KL divergence by general α -divergences remains debatable. This equivalence also implies that the asymmetry of the divergence is irrelevant to sensor management if minimization of entropy is desirable.

Our work leaves some open questions, which may be subject of further (theoretical or empirical) analysis:

- 1) Which properties the α -divergences (and also the Rényi entropies) have that would justify their use instead of the KL divergence or the Shannon entropy in sensor management?
- 2) How much of our discussion on the theoretical justification of entropy applies to multi-modal distributions which are not Gaussian mixtures?
- 3) What can be said about the meaning of entropy in distributions defined in non-Euclidean state spaces? One particularly intriguing case is the hybrid state space using by the Joint Multitarget Particle Filter [7].

VI. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the EU's Seventh Framework Programme under grant agreement n° 238710. The research has been carried out in the MC IMPULSE project: <https://mcimpulse.isy.liu.se>.

This research has been also supported by the Netherlands Organisation for Scientific Research (NWO) under the Casimir program, contract 018.003.004. Under this grant Yvo Boers holds a part-time position at the Department of Applied Mathematics at the University of Twente.

The authors would also like to thank Hans Driessen (Thales Nederland B.V.) and the reviewers for the valuable suggestions and comments.

REFERENCES

- [1] K. J. Hintz and E. S. McVey, "Multi-process constrained estimation," *IEEE Trans. Syst., Man, Cybern. A*, vol. 21, no. 1, pp. 237–244, 1991.
- [2] K. J. Hintz, "A measure of information gain attributable to cueing," *IEEE Trans. Syst. Sci. Cybern.*, vol. 21, no. 2, pp. 434–442, 1991.
- [3] J. Manyika and H. Durrant-Whyte, *Data Fusion and Sensor Management: a decentralized information-theoretic approach*. New York, NY: Ellis Horwood, 1994.
- [4] W. Schmaedeke and K. Kastella, "Event-averaged maximum likelihood estimation and information-based sensor management," in *Proc. SPIE*, vol. 2232, 1994, pp. 91–96.
- [5] K. Kastella, "Discrimination gain to optimize classification," *IEEE Trans. Syst., Man, Cybern. A*, vol. 27, no. 1, pp. 112–116, 1997.
- [6] R. Mahler, "Global optimal sensor allocation," in *Proc. Ninth National Symposium on Sensor Fusion*, Monterey, CA, 1996, pp. 167–172.
- [7] A. Doucet, B. Vo, C. Andrieu, and M. D. Thomas, "Particle filtering for multi-target tracking and sensor management," in *Proc. 5th International Conference on Information Fusion*, Annapolis, MD, 2002, pp. 474–481.
- [8] C. Kreucher, K. Kastella, and A. O. Hero, "Multi-target sensor management using alpha-divergence measures," in *Lecture Notes in Computer Science, Proc. of the 2nd International Conference on Information Processing in Sensor Networks*, no. 2634, 2003, pp. 209–222.
- [9] —, "Information based sensor management for multitarget tracking," in *Proc. SPIE*, vol. 5204, Aug. 3–8, 2003, pp. 480–489.
- [10] A. O. Hero, B. Ma, O. Michel, and J. D. Gorman, "Applications of entropic spanning graphs," *IEEE Signal Process. Mag.*, vol. 19, no. 5, pp. 85–95, 2002.
- [11] —, "Alpha-divergence for classification, indexing and retrieval," Univ. Michigan, Ann Arbor, MI, Tech. Rep. 328, May 2001.
- [12] J. M. Aughenbaugh and B. R. La Cour, "Metric selection for information theoretic sensor management," in *Proc. 11th International Conference on Information Fusion*, Cologne, Germany, Jun. 30, Jul. 1–3, 2008.
- [13] C. Kreucher, A. O. Hero, and K. Kastella, "A comparison of task driven and information driven sensor management for target tracking," in *Proc. IEEE 44th Conference on Decision and Control*, Dec. 12–15, 2005, pp. 4004–4009.
- [14] A. O. Hero, C. Kreucher, and D. Blatt, "Information theoretic approaches to sensor management," in *Foundations and Applications of Sensor Management*, A. O. Hero, D. Castañón, D. Cochran, and K. Kastella, Eds. Springer, 2007, ch. 3, pp. 33–57.
- [15] E. H. Aoki, A. Bagchi, P. Mandal, and Y. Boers, "On the "near-universal proxy" argument for theoretical justification of information-driven sensor management," in *Proc. IEEE Workshop on Statistical Signal Processing*, Nice, France, Jun. 28–30, 2011.
- [16] J. M. Aughenbaugh and B. R. La Cour, "Sensor management for particle filter tracking," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 47, no. 1, pp. 503–523, 2011.
- [17] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley–Interscience, 1991.
- [18] J. L. Williams, "Information theoretic sensor management," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, Feb. 2007.
- [19] B. Ristic, B.-N. Vo, and D. Clark, "A note on the reward function for PHD filters with sensor control," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 47, no. 2, pp. 1521–1529, 2011.
- [20] Y. Boers, J. N. Driessen, A. Bagchi, and P. Mandal, "Particle filter based entropy," in *Proc. 13th International Conference on Information Fusion*, Edinburgh, United Kingdom, Jul. 26–29, 2010.
- [21] M. Huber, T. Bailey, H. Durrant-Whyte, and U. Hanebeck, "On entropy approximation for gaussian mixture random vectors," in *Proc. IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, Aug. 20–22, 2008, pp. 181–188.
- [22] Y. Boers, E. Sviestins, and J. N. Driessen, "Mixed labelling in multitarget particle filtering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 46, no. 2, pp. 792–802, 2010.
- [23] D. Schuhmacher, B.-T. Vo, , and B.-N. Vo, "A consistent metric for performance evaluation of multi-object filters," *IEEE Trans. Signal Process.*, vol. 56, no. 8, pp. 3447–3457, 2008.
- [24] A. Rényi, "On measures of information and entropy," in *Proc. 4th Berkeley Symposium on Mathematics, Statistics and Probability*, vol. 1, 1961, pp. 547–561.
- [25] T. Hanselmann, M. Morelande, B. Moran, and P. Sarunic, "Sensor scheduling for multiple target tracking and detection using passive measurements," in *Proc. 11th International Conference on Information Fusion*, Cologne, Germany, 2008.
- [26] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Process. Mag.*, vol. 19, no. 2, pp. 61–72, 2002.
- [27] C. Kreucher, K. Kastella, and A. O. Hero, "Sensor management using an active sensing approach," *Signal Processing*, vol. 85, no. 3, pp. 607–624, Mar. 2005.
- [28] F. Liese and I. Vajda, "On divergences and informations in statistics and information theory," *IEEE Trans. Inf. Theory*, vol. 52, no. 10, pp. 4394–4412, 2006.