

ON THE “NEAR-UNIVERSAL PROXY” ARGUMENT FOR THEORETICAL JUSTIFICATION OF INFORMATION-DRIVEN SENSOR MANAGEMENT

E. H. Aoki, A. Bagchi, P. Mandal

Y. Boers

University of Twente
Department of Applied Mathematics
Enschede, The Netherlands

Thales Nederland B. V.
Hengelo, The Netherlands

ABSTRACT

In sensor management applications, sometimes it may be difficult to find a goal function that meaningfully represents the desired qualities of the estimate, such as when we do not have a clear performance metric or when the computation cost of the goal function is prohibitive. An alternative is to use goal functions based on information theory, such as the Rényi divergence (also called α -divergence).

One strong argument in favor of information-driven sensor management is that the Rényi divergence is a “near-universal” proxy for arbitrary task-driven risk functions, implying that these could be replaced by a Rényi divergence-based criterion, and this would usually result in satisfactory performance. In this paper, we present a rebuttal to that argument, which implies that finding theoretical justification for information-driven sensor management still seems to be an open problem.

Index Terms— Sensor management, information theory, Rényi divergence

1. 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 by 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).

The first suggestion of using information theory in a problem related to sensor management and state estimation

appears to have been by Hintz and McVey [1, 2], who consider Shannon entropy in target tracking using Kalman Filters. Subsequently, Manyika and Durrant-Whyte [3] consider expected information gain in sensor management and data fusion problems, and Schmaedeke and Kastella [4], Kastella [5], and Mahler [6] have considered the Kullback-Leibler (KL) divergence in various sensor management scenarios. The idea of using the more general Rényi divergence (or α -divergence), instead of the KL divergence, was introduced by Kreucher, Kastella and Hero [7].

Two subsequent works [8, 9] of Kreucher, Kastella and Hero are particularly relevant as they 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 work is organized as follows. Section 2 briefly describes the sensor management problem, and also describes the task-driven and information-driven approaches. Section 3 describes the Rényi divergence and the theoretical justification presented until so far for its use in sensor management, including the “near-universal” proxy argument. Section 4 presents a rebuttal to the “near-universal” proxy argument, and Section 5 draws the conclusions of this work.

2. THE SENSOR MANAGEMENT PROBLEM

Let us 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 $\gamma(X, Y, U)$.

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.

2.1. 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 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 [8]). 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 [8], 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.

2.2. 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 on 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.

3. THE RÉNYI DIVERGENCE AND THE “NEAR-UNIVERSAL PROXY” ARGUMENT

The Rényi divergence (or α -divergence) is a measure of difference between two distributions. Consider a pair of distributions μ and ν which admit densities $p(x)$ and $q(x)$ respectively with respect to a dominating σ -finite measure ρ . The α -divergence from μ to ν (or alternatively, from p to q) is given by

$$D_\alpha(p||q) \triangleq \frac{1}{\alpha - 1} \log \int p^\alpha(x) q^{1-\alpha}(x) \rho(dx). \quad (4)$$

where α is the order of entropy, and we apply the conventions $p^\alpha(x) q^{1-\alpha}(x) = 0$ for $p(x) = q(x) = 0$, and $a/0 = \infty$ for $a > 0$. D_0 and D_1 are defined using the limits from right and left respectively, which makes [10] D_1 the same as the Kullback-Leibler divergence.

From our previous discussion, in sensor management we would be interested in making a decision that maximizes the information gain, i.e. the difference between p_X and $p_{X|Y,U}$. Thus we can maximize the expectation of either $D_\alpha(p_{X|Y,U} || p_X)$ or $D_\alpha(p_X || p_{X|Y,U})$.

For sensor management purposes, it may seem intuitive to choose a sensing action that maximizes the “information gain” obtained by moving from prior to posterior density. The problem is to identify what this information gain means for practical purposes.

It was suggested in [11] that we use the α -divergence which, in some sense, maximizes the discrimination capability between prior and posterior distributions, which would

correspond to an optimal notion of information gain. However, through a few examples, Aughenbaugh and La Cour [12] showed that using divergences with higher discrimination capability does not necessarily lead to more informative posteriors.

3.1. The “near-universal” proxy argument

The “near-universal” proxy argument [8] is perhaps the strongest theoretical justification until so far for the use of information-driven sensor management, as it directly links it to task-driven sensor management. Let us consider some control decision u , a realization y of observation Y , and a task-driven goal $\gamma(X, y, u) = \epsilon(X, \hat{X})$. The argument claims that the expectation of γ w.r.t. $\mu_{X|Y,U}$ admits the following “sandwich-inequality”:

$$\begin{aligned} & w \exp\left(-\frac{1-\alpha_2}{\alpha_2} D_{\alpha_2}(\mu_{X|Y,U} \parallel \mu_{X|U})\right) \\ & \leq E_{\mu_{X|Y,U}}[\gamma(X, y, u)] \leq \\ & W \exp\left(-\frac{1-\alpha_1}{\alpha_1} D_{\alpha_1}(\mu_{X|Y,U} \parallel \mu_{X|U})\right) \end{aligned} \quad (5)$$

where w and W are respectively the lower bound and upper bound of γ , and the distribution $\mu_{X|U}$ may be replaced by μ_X under a few additional assumptions. There is a contradiction between [8] and [9] w.r.t. the values of α_1 and α_2 ; we will show that [9] has the correct values, i.e. $\alpha_1 > 1$ and $\alpha_2 \in [0, 1)$.

According to [9], this “sandwich-inequality” would imply that the Rényi divergence is a “near-universal” proxy that performs nearly as well as task-specific optimal policies for a wide range of tasks. By maximizing an α -divergence with $\alpha < 1$, we would automatically minimize a lower bound on the value of arbitrary error functions. Therefore we could replace possibly complex and computationally expensive task-driven risk functions by a criterion based on α -divergences and usually expect a satisfactory performance.

4. REBUTTAL OF THE “NEAR-UNIVERSAL” PROXY ARGUMENT

The “sandwich-inequality” (5) was used to justify the use of the Rényi divergence in a series of other works [13, 9, 14, 15]. Since the “near-universal” proxy argument seems very strong given the properties of information-driven measures that we have been able to identify until so far, we will attempt to repeat the same derivations done in [8] and check the validity of these results.

Let’s consider again some control decision u and a realization y of observation Y . Then let $\gamma(X, y, u)$ be an arbitrary non-negative risk function that depends on the system state X or some appropriate subset of it. We assume that γ has at least either an upper bound $W(y, u) < \infty$ or a lower bound

$w(y, u) > 0$ (for the sake of rigorousness we show explicitly the dependence of the bounds on y and u). Note that this assumption can be already quite restrictive; for instance, if $p_{X|Y,U}$ is normal and the risk function is the sum of the RMS errors, we don’t have either $W(y, u) < \infty$ or $w(y, u) > 0$.

Now, we make the assumption that $\mu_X = \mu_{X|U}$, and by noting that the support of $p_{X|Y,U}$ is contained in the support of p_X , we have

$$\begin{aligned} & E_{\mu_{X|Y,U}}[\gamma(X, y, u)] \\ & = \int \gamma(x) p_{X|Y,U}(x|y, u) \rho(dx) \\ & = \int \gamma(x) \frac{p_{X|Y,U}(x|y, u)}{p_X(x)} p_X(x) \rho(dx) \\ & = E_{\mu_X} \left[\gamma(x) \frac{p_{X|Y,U}(x|y, u)}{p_X(x)} \right]. \end{aligned} \quad (6)$$

Let’s assume that $\alpha \in [0, 1)$. For any $a > 0$, then it is easy to see that a^α is a concave function, and by applying Jensen’s inequality to (6), we have

$$\begin{aligned} & E_{\mu_{X|Y,U}}[\gamma(X, y, u)] \\ & \geq \left(E_{\mu_X} \left[\gamma^\alpha(x) \frac{p_{X|Y,U}^\alpha(x|y, u)}{p_X^\alpha(x)} \right] \right)^{\frac{1}{\alpha}} \\ & \geq \left(\int \gamma^\alpha(x) p_{X|Y,U}^\alpha(x|y, u) p_X^{1-\alpha}(x) \rho(dx) \right)^{\frac{1}{\alpha}} \\ & \geq w(y, u) \left(\int p_{X|Y,U}^\alpha(x|y, u) p_X^{1-\alpha}(x) \rho(dx) \right)^{\frac{1}{\alpha}} \\ & \geq w(y, u) \exp\left(\frac{1}{\alpha}\right) \\ & \quad \times \log \int p_{X|Y,U}^\alpha(x|y, u) p_X^{1-\alpha}(x) \rho(dx) \\ & \geq w(y, u) \exp\left(\frac{\alpha-1}{\alpha} D_\alpha(p_{X|Y,U} \parallel p_X)\right). \end{aligned} \quad (7)$$

However, because $D_\alpha \geq 0$ and $\frac{\alpha-1}{\alpha} < 0$, we have

$$w(y, u) \exp\left(\frac{\alpha-1}{\alpha} D_\alpha(p_{X|Y,U} \parallel p_X)\right) \leq w(y, u) \quad (8)$$

and because of the trivial relation

$$E_{\mu_{X|Y,U}}[\gamma(X, y, u)] \geq w(y, u) \quad (9)$$

we observe that the bound (7) has no real significance, since regardless of the value of the Rényi divergence, it is always more loose than the trivial lower bound $w(y, u)$, that do not predictably depend on the Rényi divergence. An analogous problem happens for $\alpha > 1$.

So, it becomes evident that the described bounds do not give any performance guarantees with respect to using α -divergences as replacement for task-driven sensor management criterions, and thus we cannot consider these bounds as a theoretical justification.

5. CONCLUSIONS

In this paper we reviewed existing studies on theoretical justification of the Rényi divergence as sensor management criterion, and presented a refutation to the “near-universal” proxy argument. This implies that, to the best of our knowledge, there is no such strong relationship between task-driven and information-driven sensor management. Therefore, although the use of information-theoretic divergences in sensor management seem to have been validated by various empirical studies, from a theoretical point of view it is still not completely clear how these criteria would lead to generation of good quality estimates.

We should make clear that the purpose of our work is not to disqualify the use of information-driven sensor management, neither to neglect the advances made in the field by previous authors. On the contrary, we recognize that all previous results, both theoretical and empirical, have greatly contributed to a better comprehension of information-theoretic measures and their relationship with the estimation and sensor management problems. Consequently, we could say that the purpose of this work was mostly to motivate the community by pointing out that there is still work to be done.

6. 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.) for contributing to the paper.

7. REFERENCES

- [1] K. J. Hintz and E. S. McVey, “Multi-process constrained stimulation,” *IEEE Trans. Syst., Man, Cybern.*, 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] 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, Apr. 22–23, 2003, pp. 209–222.
- [8] 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.
- [9] 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.
- [10] T. van Erven and P. Harremoës, “Rényi divergence and majorization,” in *Proc. IEEE International Symposium on Information Theory*, Austin, TX, Jun. 13–18, 2010, pp. 1335–1339.
- [11] C. Kreucher, K. Kastella, and A. O. Hero, “Information based sensor management for multitarget tracking,” in *Proc. SPIE*, vol. 5204, Aug. 3–8, 2003, pp. 480–489.
- [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, K. Kastella, and M. R. Morelande, “An information-based approach to sensor management in large dynamic networks,” in *Proc. IEEE*, vol. 95, no. 5, May 2007, pp. 978–999.
- [14] A. O. Hero and C. Kreucher, “Network sensor management for tracking and localization,” in *Proc. 10th International Conference on Information Fusion*, Québec, Canada, Jul. 9–12, 2007.
- [15] E. K. P. Chong, C. Kreucher, and A. O. Hero, “Partially observable Markov decision process approximations for adaptive sensing,” *Discrete Event Dynamic Systems*, vol. 19, no. 3, pp. 377–422, Sep. 2009.