

AN INFORMATION THEORETIC APPROACH TO PROCESSING MANAGEMENT

Chris Kreucher and Kevin Carter

The University of Michigan
Ann Arbor, MI

ABSTRACT

In region surveillance applications, sensors oftentimes accumulate an overwhelmingly large amount of data, making it infeasible to process all of the collected data in real-time. For example, a multi-channel synthetic aperture radar (SAR) flown on an airborne platform could receive on the order of 10 GBits of data per second. This data can be exploited in a number of ways (e.g., constructing a detected image, applying an ATR algorithm, or performing moving target processing) each of which requires significant computational resources. Given the enormous amount of data and the correspondingly large number of potential exploitation algorithms, there simply are not enough computational resources to process all of the data with all possible exploitation algorithms.

The natural question then becomes one of how to most effectively utilize limited processing resources so as to facilitate real time exploitation of the collected data. This paper presents an information theoretic approach for processing action selection which is predicated on predicting the amount of information flow each potential processing action is expected to generate. The aim is to select those exploitation algorithms (and, in general, the physical region and algorithm parameter settings) that will be most useful in refining the underlying estimate of the surveillance region state. We show by simulation on a model problem that the information theoretic method is able to outperform other methods of processing selection.

Index Terms— resource management, information theory, particle filtering, joint multitarget probability density, multitarget tracking

1. INTRODUCTION

This paper is concerned with the problem of using limited processing resources most effectively. The method we propose is a novel combination of predictive density estimation and information theoretic optimization which predicts the amount of information that is expected to be gained for each candidate method of resource utilization.

We focus on a region surveillance application, where an airborne sensor (or collection of airborne sensors) is charged with detecting, tracking, and classifying moving targets. Each

sensor collects an enormous amount of data and the desire is to most effectively process that data to determine the contents of the surveillance region. Due to the sheer amount of data collected, it is infeasible to simply process all of the data through all of the available exploitation algorithms. Therefore, a method of selecting which part of the collected data to process, which processing algorithm to use, and what parameter settings are best is required.

In this paper, we propose a method of processing management (i.e., selecting how to use the available processors) based on predicting the amount of information flow that will result from each candidate processing action and selecting the best. Information flow is a nice metric for a number of reasons. First, it ably balances the desire to sharpen ones estimate about the number of targets with the desire to sharpen estimates about the kinematic states (i.e., position and velocity) and the classification of each target [1]. Second, information theoretic methods have been shown to bound any risk based criteria, and hence they provide a universal metric [2]. Other relevant work focuses on the dual problem of sensor management, e.g., [3, 4].

This paper proceeds as follows. Section 2 provides a cursory overview of Bayesian filtering for target state estimation. The main idea is to construct a probability density which describes the state of a moving target by synthesizing sensor measurements, sensor models, and target models optimally. Section 3 provides the new work reported in this paper. The principle is to use information theory to predict which processing action will be the most useful, and to then select that processing action from the host of candidate actions. Section 4 provides a simple simulation result on a model problem that illustrates the efficacy of the proposed method.

2. BAYESIAN FILTERING FOR STATE ESTIMATION

For simplicity of exposition, we consider here the single target tracking problem, where it is known *a priori* there is one target and furthermore the initial state is also known. These assumptions are made for notational simplicity and can be relaxed. See, for example, [5] where the joint multitarget detection and tracking problem is treated fully using the same methods given here.

In the Bayesian approach, one constructs the probability

The authors can be reached at {ckreuche,kmcarter}@umich.edu

density $p(x^k|Z^k)$ recursively. This describes the probability a target is in state x at time k given the set of all observations made up to and including time k , denoted Z^k . The target state x may, for example, consist of a two dimensional position and velocity, i.e., $x = [x, x', y, y']$. We leave this general for the time being. The observation history Z^k is a combination of all observations made previously, i.e., $Z^k = \bigcup_{i=1 \dots k} z_i$, where each z_i can be a scalar, vector, or matrix. In our context, an observation z_k occurs when a decision is made to use the data received at time k with a particular exploitation algorithm r .

Notionally, the probability density is constructed recursively by first applying the Chapman-Kolmogorov temporal update (i.e., prediction):

$$p(x^k|Z^{k-1}) = \int dx p(x^k|x^{k-1})p(x^{k-1}|Z^{k-1}) , \quad (1)$$

and then the measurement update using Bayes' rule:

$$p(x^k|Z^k) = \frac{p(x^k|Z^{k-1})p(z^k|x^k)}{p(z^k|Z^k)} . \quad (2)$$

This recursion is based on the target motion model $p(x^k|x^{k-1})$ which describes probabilistically the state of a target at a future time conditioned on its current state, and the sensor model $p(z^k|x^k)$ which describes the coupling between a sensor's measurements and the unknown state.

In the case where the state transition density $p(x^k|x^{k-1})$ and the measurement likelihood $p(z^k|x^k)$ are of a particular form (linear and Gaussian), these equations yield the commonly used Kalman Filter recursions. In this case, the probability density can be represented exactly by its mean and covariance. In the more general case, an alternative method is required. In this work, we choose to represent the (potentially non-Gaussian) state probability density using a particle filter. The particle filter representation says the probability density is approximated by a set of N weighted samples, i.e.,

$$p(x|Z) \approx \sum_{p=1}^N w_p \delta(x - x_p) . \quad (3)$$

To do this tractably (especially in the multitarget case), one requires a sophisticated approach to constructing the importance (sampling) density. The details of this approach are omitted here, but can be found elsewhere [5].

3. INFORMATION THEORY FOR PROCESSING MANAGEMENT

Construction of the single target PDF yields a probability density (represented by a set of weighted samples) on the target state x conditioned on all observations made. Conceptually speaking, if this probability density has high entropy, it represents great uncertainty in the target state. Analogously, a low entropy density implies high level of certainty about the

target state. It is for this reason, that we suggest a method of processing management based on information theory.

The main idea of the information theoretic approach to processing management is as follows. Given a prediction of the target state $p(x^k|Z^{k-1})$ and a model on how the sensor works $p(z^k|x^k)$, we predict (in advance) what processing action r will yield the maximum benefit, where benefit is measured in terms of information flow. This is done by measuring the information flow between the prior density $p(x^k|Z^{k-1})$ and the posterior $p(x^k|Z^k)$. Large information flow indicates that the new observation has added significant information (i.e., significantly reduced uncertainty). Of course, one cannot compute the amount of information flow until after the processing action is taken and the actual observation z^k is made. We therefore propose to compute the *expected* amount of information flow that would result if a particular processing decision was made and use this to decide which processing action to take.

3.1. The Rényi Divergence

In our approach, the calculation of information gain between two densities p_1 and p_0 is done using the Rényi information divergence [1], also known as the α -divergence:

$$D_\alpha(p_1||p_0) = \frac{1}{\alpha - 1} \ln \int p_1^\alpha(x)p_0^{1-\alpha}(x)dx . \quad (4)$$

The function D_α in eq. (4) is a measure of the divergence between the densities p_0 and p_1 . In the present application, we wish to compute the divergence between the prediction density $p(x^k|Z^{k-1})$ and the updated density after the observation z^k is made when performing processing action r^k , denoted $p(x^k|Z^{k-1}, z^k, r^k)$. Notice that we now include the processing action taken at time k , r^k , explicitly into the notation for clarity. This divergence measures the amount of information that the new observation has provided and allows us to rank the utility of different processing decisions according to the information flow they produce. The relevant divergence for our setting is thus given by

$$D_\alpha \left(p(\cdot|Z^{k-1}, z^k, r^k) || p(\cdot|Z^{k-1}) \right) = \frac{1}{\alpha - 1} \times \ln \int p^\alpha(x^k|Z^{k-1}, z^k, r^k) p^{1-\alpha}(x^k|Z^{k-1}) dx^k . \quad (5)$$

Using Bayes' formula (eq. (1)), we obtain

$$D_\alpha \left(p(\cdot|Z^{k-1}, r^k, z^k) || p(\cdot|Z^{k-1}) \right) = \frac{1}{\alpha - 1} \times \ln \frac{1}{p^\alpha(z^k|Z^{k-1}, r^k)} \int p^\alpha(z^k|x^k, r^k) p(x^k|Z^{k-1}) dx^k , \quad (6)$$

which shows that the ingredients to computing the divergence are the prediction density $p(x^k|Z^{k-1})$, the measurement likelihood $p(z^k|x^k, r^k)$ and the received observations z^k . Under

the particle filter approximation to the posterior, this integral becomes the discrete sum

$$D_\alpha \left(p(\cdot | Z^{k-1}, r^k, z^k) || p(\cdot | Z^{k-1}) \right) \approx \frac{1}{\alpha - 1} \ln \frac{1}{\left(\sum_{p=1}^N w_p p(z^k | x_p) \right)^\alpha} \sum_{p=1}^N w_p p^\alpha(z^k | x_p) . \quad (7)$$

3.2. Expected Rényi Divergence for a Processing Action

To determine the best action to take next, we must in fact predict the value of taking action r^k *before actually making* the observation z^k . Therefore, we calculate the *expected value* of the divergence for each possible action and use this to select the next action. The expectation may be written as an integral over all possible outcomes z^k when taking action r^k as

$$\mathbb{E} \left[D_\alpha \left(p(\cdot | Z^{k-1}, z^k, r^k) || p(\cdot | Z^{k-1}) \right) | Z^{k-1}, r^k \right] = \int dz^k p(z^k | Z^{k-1}) D_\alpha \left(p(\cdot | Z^{k-1}, z^k, r^k) || p(\cdot | Z^{k-1}) \right) . \quad (8)$$

The expectation is across the observation value z^k and is to be interpreted as a conditional expectation where the past observations Z^{k-1} , past sensor actions, and current sensing action r^k are known.

Then the method of scheduling we advocate is to choose the best action \hat{r}^k as the one that maximizes the expected gain in information, i.e.,

$$\hat{r}^k = \arg \max_{r^k} \mathbb{E} \left[D_\alpha \left(p(\cdot | Z^{k-1}, z^k, r^k) || p(\cdot | Z^{k-1}) \right) | Z^{k-1}, r^k \right] . \quad (9)$$

Under the assumption that z is thresholded (i.e., $z = 0$ or $z = 1$ for a detection and non-detection, respectively), and using the particle filter approximation, this expectation becomes a simple discrete sum.

4. SIMULATION RESULTS

This section presents a simulation result illustrating the information theoretic method of processing management.

We consider the model problem illustrated in Figures 1 and 2. There is a single airborne sensor charged with tracking a moving ground target. The sensor has two “modes”: a ground moving target indication (GMTI) mode, and a synthetic aperture radar (SAR) mode.

In GMTI mode, a target that is moving with sufficient radial velocity with respect to the sensor is detectable with probability P_d^{GMTI} . The sensor resolution (pixel size) is given by dx^{GMTI} and dy^{GMTI} . Each pixel has a false alarm probability of P_f^{GMTI} . This sensor model simulates the processing steps done to construct an MTI image and extract the movers.

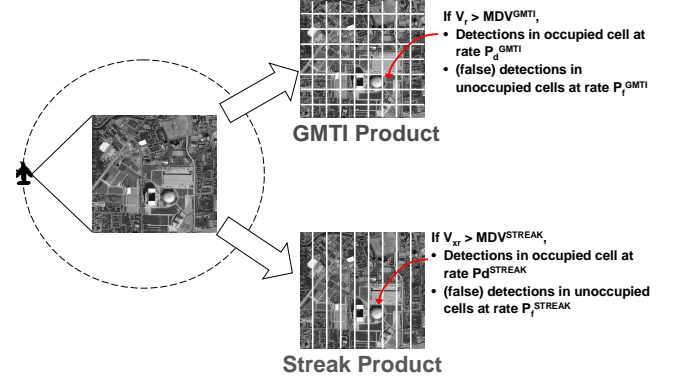


Fig. 1. The simulation setup. An airborne platform is to track a moving target. Due to computational limitations, it must choose which exploitation algorithm to use at each time step.

In SAR mode, a target moving with sufficient cross radial velocity with respect to the sensor is detectable using “streak-detection” with probability P_d^{STREAK} . The sensor resolution (pixel size) is given by dx^{STREAK} and dy^{STREAK} , where $dy^{STREAK} \gg dy^{GMTI}$. Each pixel has a false alarm probability of P_f^{STREAK} . This sensor model simulates the processing steps done to construct a SAR image and extract movers using a technique known as streak detection.

The initial target position is known, but its subsequent motion is not. The target moves along a trajectory that at some points has it moving radially away from the sensor and at other points cross-radially, as shown in Figure 2.




Fig. 2. The trajectory of the moving target. The target moves both radially and cross radially at different times leading to the need to intelligently select the sensing modality.

We are interested in a method of processing management that most effectively makes use of the limited ability to process data. It is assumed for the purposes of this simu-

lation that it is computationally infeasible to simply process the collected data with both methods, and hence we need to select whether to use SAR or GMTI at each time step of the algorithm (on the fly). We compare the performance of the following methods:

- A method that always uses SAR
- A method that always uses GMTI
- A rule-based method, which computes the MMSE estimate of target velocity and selects SAR or GMTI based on radial and cross radial velocity estimates.
- The information theoretic method, which computes the expected gain in information for SAR and GMTI and selects the one that is expected to maximize information flow.
- A method which uses both the SAR and GMTI modes (this approach, assumed infeasible, provides an upper bound of performance).

Figure 3 shows the performance of each of the methods, as measured by the ability to successfully track a single moving target. As expected, the (assumed infeasible) method that uses both processing algorithms provides the best performance. Correspondingly, the methods that simply use one mode exclusively perform very poorly. The SAR method loses the target when its cross radial velocity is low while the GMTI method loses the target when its radial velocity is low. The rule-based method is able to adaptively switch between the two algorithms and (typically) use GMTI when the target is projected to be moving radially and SAR otherwise. The information-based method adds an additional level of sophistication in that it explicitly incorporates filter uncertainty about target velocity into the picture, resulting in an enhanced ability to predict which method is appropriate. The resulting performance is thus superior to the rule-based method.

5. CONCLUSION

This paper has described a method of selecting processing actions based on information theory. Simulation results show the potential power of this approach.

There are a number of avenues for future work. Many of these have at their heart a potential computational explosion. First, there is the $O(M \text{ choose } N)$ problem of scheduling N processors, where the goal is to select N processing actions from a set of M possibilities ($M \gg N$) jointly. A related problem occurs when the possible processing actions are not enumerable, e.g., when an algorithm parameter drawn from the continuum is to be selected. Furthermore, a method of accounting for the potential different amounts of time each processing action takes (e.g., by optimizing information rate

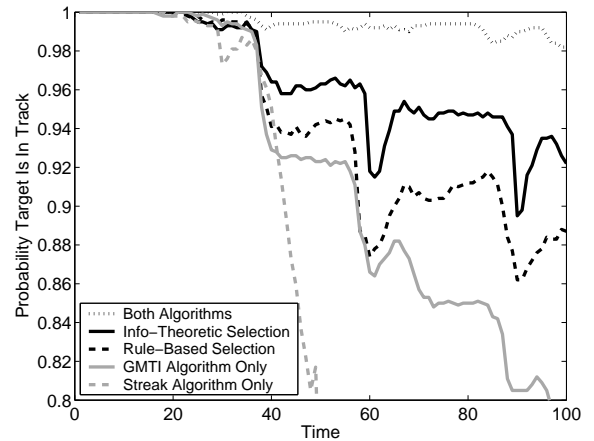


Fig. 3. The information theoretic method of processing action selection outperforms other methods of choosing processing mode. The upper bound of performance (achieved using both processing modes at each time) is shown for comparison.

rather than simply information flow) is required. This question belongs in the more general class of extensions of the approach to multi-step optimization (instead of greedily choosing the next best action).

6. REFERENCES

- [1] C. Kreucher, A. O. Hero, K.D. Kastella, and M.R. Morelande, "An information-based approach to sensor management in large dynamic networks," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 978–999, May 2007.
- [2] C. Kreucher, A. Hero, and K. Kastella, "A comparison of task driven and information driven sensor management for target tracking," in *The 44th IEEE Conference on Decision and Control*, December 12-15 2005, pp. 4004–4009.
- [3] F. Zhao, J. Shin, and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Processing Magazine*, pp. 61–72, March 2002.
- [4] J. L. Williams, J. W. Fisher, and A. S. Willksy, "Approximate dynamic programming for communication-constrained sensor network management," *IEEE Transactions on Signal Processing*, vol. 55, no. 8, August 2007.
- [5] C. Kreucher, K. Kastella, and A. Hero, "Multitarget tracking using the joint multitarget probability density," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, no. 4, pp. 1396–1414, October 2005.