

Variance Reduction for Monte Carlo Implementation of Adaptive Sensor Management *

Sumeetpal Singh[†], Ba-Ngu Vo[‡], Robin J. Evans[‡] and Arnaud Doucet[†]

[†]Signal Processing Group, Engineering Department
University of Cambridge, Cambridge CB2 1PZ, UK

[‡]Department of Electrical & Electronic Engineering,
The University of Melbourne, VIC 3010, Australia

sss40@eng.cam.ac.uk, b.vo@ee.mu.oz.au, r.evans@ee.mu.oz.au, ad2@eng.cam.ac.uk

Abstract – Adaptive sensor management (scheduling) is usually formulated as a finite horizon POMDP and implemented using sequential Monte Carlo. In Monte Carlo, variance reduction is important for the reliable performance of the sensor scheduler. In this paper, we propose a Control Variate method for variance reduction when the sensor is scheduled using the Kullbach Leibler criterion.

Keywords: Tracking, sensor management, filtering, particle filter, variance reduction, control variate.

1 Introduction

Consider a target whose state $X_k \in \mathbf{R}^{n_x}$ evolves in a Markov fashion,

$$X_k \sim p_k(\bullet | X_{k-1}) \quad (1)$$

with initial density π_0 , where $k = 1, 2, \dots$ denotes discrete-time (The notation “ \sim ” implies distributed according to). To track this target, a finite set \mathbb{A} of allocatable sensors are available. The observation $Y_k \in \mathbf{R}^{n_y}$ received by sensor $A_k \in \mathbb{A}$ is described by

$$Y_k \sim q_k(\bullet | X_k, A_k). \quad (2)$$

Suppose that at each time k , only measurements from a single sensor may be received due to limited communication bandwidth. The problem of scheduling the sequence of sensors to optimise the tracking performance is known as *sensor management or scheduling*.

*This research is supported by DSTO (CENDS), ARC and DARPA under ONR Grant Number N00014-02-1-0802. This work has been carried out when the first author was working at the University of Melbourne

Adaptive sensor scheduling can be formulated as a finite horizon partially observed Markov decision process (POMDP). For a POMDP, the *filtering density* is a sufficient statistic and the optimal scheduler (or *policy*) can be constructed using Dynamic Programming (DP) [1]. In this paper, we only consider a *one-step look ahead* scheduler where at each time step, the sensor that maximises the expected differential entropy (or Kullbach Leibler discrimination) between the filtering and the predicted density is selected. The use of the Kullbach Leibler criterion is motivated in [4]. Effectively, the sensor that brings the most information about the target being tracked is chosen.

The design of a sensor scheduler, among other issues, hinges on the ability to calculate and propagate the filtering density in time. In [4], a sensor scheduler based on the Kullbach Leibler (KL) discrimination criterion was implemented using Sequential Monte Carlo (SMC). SMC methods are powerful tools for approximating the filtering density of interest even when the systems has non-linear non-Gaussian dynamics. Additionally, SMC is an asymptotically consistent (in the number of particles) approximation scheme that is very simple to implement unlike the extended Kalman filter, deterministic integration methods, etc. However, as noted in [4], SMC methods suffer from high variance and can deteriorate the performance of the scheduler. For example, if the SMC estimate of the objective function for the various sensors have overlapping confidence regions, then there is a high probability of selecting the wrong sensor.

Variance reduction is crucial in SMC based sensor management. An obvious way to reduce the variance of the SMC estimate is to increase the number of sam-

ples (or particles), but this is computationally inefficient. In this paper, we present a Control Variate (CV) method for variance reduction when the KL criterion is used to discriminate between sensors. The proposed method, unlike an importance sampling based variance reduction method, does not require any modification to the process of generating the samples (or particle filter). Rather, it only collects additional statistics from these samples and it is very simple to implement. We develop the proposed methodology in detail for a target tracking application and provide some asymptotic results.

2 Formulation of Sensor Management and SMC Implementation

Given the sequence of actions $a_{1:k} \equiv [a_1, \dots, a_k]$ and measurements $y_{1:k} \equiv [y_1, \dots, y_k]$, the *filtering density* and *predicted density* at time k are denoted respectively by π_k and $\pi_{k+1|k}$ (or $\pi_k^{(y_{1:k}, a_{1:k})}$ and $\pi_{k+1|k}^{(y_{1:k}, a_{1:k})}$) to emphasise the dependence on $(y_{1:k}, a_{1:k})$. For any $y, a, k > 0$, the *prediction* and *update operators* $\Psi_k^{(a)}$ and $\Upsilon_k^{(y,a)}$ are defined respectively for any $f \in \mathcal{P}(\mathbf{R}^{n_x})$, the space of all densities on \mathbf{R}^{n_x} , by

$$\Psi_k^{(a)}[f](x) := \int p_k(x|x_{k-1}, a) f(x_{k-1}) dx_{k-1} \quad (3)$$

$$\Upsilon_k^{(y,a)}[f](x) := \frac{f(x)q_k(y|x, a)}{\langle f, q_k(y|\cdot, a) \rangle} \quad (4)$$

where $\langle f, g \rangle \equiv \int f(x)g(x)dx$. (Not to be confused with the inner product defined in Section 3.) The prediction and update operators have the subscript k as we assume an inhomogeneous model. Additionally, the a -dependence in (3) is only to point out that the framework is general enough to allow a target model where the target manoeuvre can depend on the choice of the active sensor. Such a model is relevant for active sensors that can be detected by the target. The Bayes recursion is given by

$$\pi_{k|k-1} = \Psi_k^{(a_k)}[\pi_{k-1}], \quad (5)$$

$$\pi_k = \Upsilon_k^{(y_k, a_k)}[\pi_{k|k-1}]. \quad (6)$$

Define the *filtering operator* $\Phi_k^{(y_k, a_k)}$ to be the product operator $\Upsilon_k^{(y_k, a_k)}\Psi_k^{(a_k)}$, i.e., $\pi_k = \Phi_k^{(y_k, a_k)}[\pi_{k-1}] = \Upsilon_k^{(y_k, a_k)}[\Psi_k^{(a_k)}[\pi_{k-1}]]$. Applying the filtering operator repeatedly yields for any $k \geq 1$,

$$\pi_k^{(y_{1:k}, a_{1:k})} = \left(\prod_{i=1}^k \Phi_i^{(y_i, a_i)} \right) [\pi_0] =: \Phi_{1:k}^{(y_{1:k}, a_{1:k})}[\pi_0] \quad (7)$$

Given the system dynamics (1-2), it is evident from (7) that the filtering density at time k is completely determined by $y_{1:k}$, $a_{1:k}$, and π_0 .

Our objective function for sensor selection is the expected differential entropy between the filtering density π_k and the predictive density $\pi_{k|k-1}$. The differential entropy or KL discrimination between two densities f and g is given by:

$$r(f, g) := \int f(x) \log \left(\frac{f(x)}{g(x)} \right) dx. \quad (8)$$

$r(f, g)$ is non-negative and $r(f, g) = 0$ if $f = g$. By selecting the sensor that maximises $r(\pi_k, \pi_{k|k-1})$, we maximise the information yielded at the k th observation epoch. The KL discrimination $r(\pi_k, \pi_{k|k-1})$ can be expressed in terms of the filtering density at $k-1$, action a_k and observation y_k as follows:

$$\begin{aligned} r(\pi_k, \pi_{k|k-1}) &= r(\Phi_k^{(y_k, a_k)}[\pi_{k-1}], \Psi_k^{(a_k)}[\pi_{k-1}]) \\ &=: h_k(\pi_{k-1}, a_k, y_k). \end{aligned} \quad (9)$$

As in [4], one averages over the possible observations at time k to yield the expected KL discrimination is

$$\begin{aligned} &J(\pi_{k-1}; a) \\ &:= \int \int h_k(\pi_{k-1}, a, y) q_k(y|x, a) \pi_{k|k-1}(x) dy dx, \\ &= \int \int \log \left(\frac{q_k(y|x, a)}{\langle q_k(y|\cdot, a), \pi_{k|k-1} \rangle} \right) \\ &\quad \times q_k(y|x, a) \pi_{k|k-1}(x) dx dy. \end{aligned} \quad (10)$$

The optimal 1-step ahead policy is given by

$$a_k^* = \arg \max_a J(\pi_{k-1}; a) \quad (11)$$

From a practical viewpoint, real tracking models are often non-linear and non-Gaussian, and the filtering density as well as the expected KL discrimination cannot be expressed in closed-form. However, these quantities can be approximated using SMC.

At time $k-1$, let $\{x_{k-1}^{(i)}, w_{k-1}^{(i)}\}_{i=1}^L$ be a particle approximation of π_{k-1} where $x_{k-1}^{(i)}$ is the sample of the state and $w_{k-1}^{(i)}$ is the positive weight. For $i = 1, \dots, L$, sample $x_k^{(i)} \sim p_k(\cdot|x_{k-1}^{(i)})$, and $y_k^{(i)} \sim q_k(\cdot|x_k^{(i)}, a)$. Then, $\{(y_k^{(i)}, x_k^{(i)}), w_{k-1}^{(i)}\}_{i=1}^L$ is a particle approximation of $q_k(y|\cdot, a)\Psi_k[\pi_{k-1}]$. The SMC approximation of $J(\pi_{k-1}; a)$ is given by

$$\hat{J}(\pi_{k-1}; a) = \sum_{i=1}^L w_{k-1}^{(i)} \log \left(\frac{q_k(y_k^{(i)}|x_k^{(i)}, a)}{\sum_{j=1}^L w_{k-1}^{(j)} q_k(y_k^{(j)}|x_k^{(j)}, a)} \right). \quad (12)$$

Advanced particle methods rely on importance densities other than $p_k(x_k|x_{k-1})$; this has been shown to improve results significantly. In the context of sensor scheduling, it is also possible to use “advanced” importance sampling techniques to estimate $J(\pi_{k-1}; a)$. However, given $\{(y_k^{(i)}, x_k^{(i)}, w_{k-1}^{(i)})\}_{i=1}^L$, our concern here is to reduce the variance in the approximation (12).

3 Variance Reduction by Control Variate (CV)

The accuracy of the SMC estimates of the objective function affects the ranking of the sensors. To see this, let δ_a denote the standard deviation of the estimate $\hat{J}(\pi_{k-1}; a)$ and consider the confidence interval $I_a := [\hat{J}(\pi_{k-1}; a) - \delta_a, \hat{J}(\pi_{k-1}; a) + \delta_a]$ for each $a \in \mathbb{A}$. Suppose that $a^{(1)}$ and $a^{(2)}$ are the 1st and 2nd rank sensors respectively. If the intervals $I_{a^{(1)}}$ and $I_{a^{(2)}}$ has a large overlap, then there is a high chance that these two sensors are incorrectly ranked. The situation is even worse if many of the intervals I_a , $a \in \mathbb{A}$ overlaps with each other. Thus, finding the optimal sensor at each time step hinges on the ability to compute the objective function accurately. In this section, we discuss the use of the control variate method to reduce the variance in the estimation of $\hat{J}(\pi_{k-1}; a)$.

3.1 Background

The method of CV is a widely used variance reduction method for Monte Carlo simulation. Recently, it has been shown that other popular variance reduction techniques such as conditional Monte Carlo, antithetics, rotation sampling, stratification can be viewed as various implementations of the CV method [2]. The CV method works as follows: consider the problem of estimating $\alpha = E(X)$ for some random variable X . If $x^{(i)}$, $i = 1, \dots, L$, are L i.i.d. samples of X , then one could use the sample mean estimator which incurs a variance of $\text{var}(X)/L$. In many situations, one may exploit the additional structure in the problem to reduce the variance of the sample mean estimator. Assume that there is a zero-mean random variable Y that is jointly distributed with X . Then, the estimator

$$L^{-1} \sum_{i=1}^L x^{(i)} + uy^{(i)}, \quad (13)$$

where $(x^{(i)}, y^{(i)})$ are i.i.d. realizations of (X, Y) , is also unbiased for any fixed scalar $u \in \mathbf{R}$. In the literature, Y and u are known as the CV and CV coefficient respectively. It is easy to show that the optimal CV coefficient that minimises the variance of the estimator

$$(13) \text{ is } u^* = -E(XY)/\text{var}(Y). \quad (14)$$

The variance of (13) at u^* is

$$L^{-1} \left(\text{var}(X) - \frac{E(XY)^2}{\text{var}(Y)} \right), \quad (15)$$

which is strictly less than $\text{var}(X)/L$ unless $E(XY) = E(X)E(Y)$. In some cases, it is possible to construct a *perfect CV*, i.e. a random variable Y for which

$$\text{var}(X + u^*Y) = 0.$$

(We give an example below.) If a perfect CV exists, then practically it means that we can have a near perfect estimator of $\alpha = E(X)$ for a given sample size L .

3.1.1 Asymptotics

The optimal CV constant in (14) needs to be estimated in implementation. Given the i.i.d. realizations $(x^{(i)}, y^{(i)})$, $i = 1, \dots, L$, one could use the estimator

$$L^{-1} \sum_{i=1}^L x^{(i)} + u_{i-1}y^{(i)} \quad (16)$$

where

$$u_0 = 0, \quad u_k := \frac{\sum_{i=1}^k x^{(i)}y^{(i)}}{\sum_{i=1}^k (y^{(i)})^2}, \quad k \geq 1. \quad (17)$$

In practice,

$$L^{-1} \sum_{i=1}^L x^{(i)} + u_L y^{(i)}, \quad (18)$$

is favourable as the final value u_L would be a better estimate of u^* . However, the asymptotics are more easily studied for (16) as $\text{cov}(x^{(i)} + u_{i-1}y^{(i)}, x^{(j)} + u_{j-1}y^{(j)}) = 0$ for $i \neq j$. Let

$$\nu(u) := \text{var}(X + uY). \quad (19)$$

It was shown in [3] that if $u_L \rightarrow u^*$ (in probability if not almost surely), then

$$\begin{aligned} & L \times \text{var} \left(L^{-1} \sum_{i=1}^L x^{(i)} + u_{i-1}y^{(i)} \right) \\ &= L^{-1} \sum_{i=1}^L E(\nu(u_{i-1})) \\ &\rightarrow \nu(u^*) \end{aligned} \quad (20)$$

as $L \rightarrow \infty$. (Sharper bounds on the convergence rate can be given in several scenarios, but these bounds are specific to the implementation [3].) We may interpret

this result as follows. In the standard sample mean estimator, i.e. (20) with $u_i = 0$ for all i , the variance is $\nu(0)/L$. In the CV version, the variance is

$$\left(L^{-2} \sum_{i=1}^L E(\nu(u_{i-1})) - L^{-1}\nu(u^*) \right) + L^{-1}\nu(u^*).$$

As the term in the brackets tends to 0 faster than $L^{-1}\nu(u^*)$, for large enough L , the variance of the estimator is approximately $L^{-1}\nu(u^*)$. If an optimal CV exists, the estimator is near perfect.

3.2 CV for Entropy Estimation

We may rewrite (10) as the difference of the entropies

$$\begin{aligned} & \int \log(q_k(y|x, a) \Psi_k[\pi_{k-1]}(x)) \\ & \times q_k(y|x, a) \Psi_k[\pi_{k-1]}(x) dx dy \\ & - \int \log(\langle q_k(y|\cdot, a), \Psi_k[\pi_{k-1}] \rangle) \\ & \times \langle q_k(y|\cdot, a), \Psi_k[\pi_{k-1}] \rangle dy. \end{aligned} \quad (21)$$

Thus, in this subsection, we discuss variance reduction for the sample mean estimator

$$L^{-1} \sum_{i=1}^L \log f(x^{(i)}), \quad x^{(i)} \sim f,$$

of the entropy $\int f(x) \log f(x) dx$ of some probability density f .

Consider a family of densities f_θ , $\theta \in \Theta$. (Examples of $\{f_\theta, \theta \in \Theta\}$ are given in Proposition 2 and Section 4.) Let $\theta_0 \in \Theta$ be a fixed parameter. Suppose that we wish to estimate the entropy $\int f_{\theta_0}(x) \log f_{\theta_0}(x) dx$. For the remainder of this section, we denote the random vector in \mathbf{R}^n with density f_{θ_0} by W . For functions $X, Y : \mathbf{R}^n \rightarrow \mathbf{R}$, let $\langle X, Y \rangle := E(X(W)Y(W))$ and $\|X\| = \langle X, X \rangle^{1/2}$. $\langle \cdot, \cdot \rangle$ is an inner product on the space of functions from \mathbf{R}^n to \mathbf{R} that have a finite second moment while $\|\cdot\|$ is the norm induced by the inner product. Define

$$X(W) := \log f_{\theta_0}(W). \quad (22)$$

The object of interest is

$$\alpha = E(X(W)) = \int f_{\theta_0}(x) \log f_{\theta_0}(x) dx, \quad (23)$$

which can be estimated by $L^{-1} \sum_{i=1}^L X(W^{(i)})$. To implement a CV version, we require a random variable $Y(W)$ satisfying $E(Y(W)) = 0$. It is straightforward to show that

$$\frac{\min_u \text{var}(X(W) + uY(W))}{\text{var}(X(W))} = 1 - \left\langle \frac{X - \alpha}{\|X - \alpha\|}, \frac{Y}{\|Y\|} \right\rangle^2. \quad (24)$$

Thus, the inner product on the right hand-side may be viewed as the variance reduction factor. Note that if the projection of $(X - \alpha)/\|X - \alpha\|$ on Y is close to 1, i.e. the angle between the 2 vectors $X - \alpha$ and Y is small, the variance reduction is large.

The idea suggested by (24) is to find a zero-mean random variable that is closely aligned to $\bar{X} := X - \alpha$, the projection of X on the zero-mean subspace. In this paper, we propose the use of the score

$$S_\theta(W) := \frac{f_\theta(W)}{f_{\theta_0}(W)} - 1, \quad (25)$$

for entropy estimation. The score satisfies

$$E(S_\theta(W)) = 0$$

and can be implemented as a CV. The questions are what is the relationship between the score S_θ and X ?, more importantly, what can we say about the alignment of the score with \bar{X} .

We now shed some light as to why the score might serve well as a CV. Define for each $\theta \in \Theta$, the random variables

$$X_\theta(W) := \log f_\theta(W). \quad (26)$$

Note that $X = X_{\theta_0}$.

Proposition 1 Assume

$\|X_\theta - X - S_\theta\| / \|f_\theta - f_{\theta_0}\| \rightarrow 0$ as $\theta \rightarrow \theta_0$. Then

$$\lim_{\theta \rightarrow \theta_0} \left\| \frac{X_\theta - X}{\|X_\theta - X\|} - \frac{S_\theta}{\|S_\theta\|} \right\| = 0$$

(All proofs appear in [7].)

(Equivalently, $\lim_{\theta \rightarrow \theta_0} \left\langle \frac{X_\theta - X}{\|X_\theta - X\|}, \frac{S_\theta}{\|S_\theta\|} \right\rangle = 1$.) The score $S_\theta(W)$ is the first order term in the Taylor expansion of $\log f_\theta(W) - \log f_{\theta_0}(W)$ about $f_{\theta_0}(W)$, i.e.,

$$\log f_\theta(W) = \log f_{\theta_0}(W) + S_\theta(W) + \text{higher order terms.}$$

The above proposition asserts that asymptotically (as θ tends to θ_0), the score S_θ can be treated as the difference $X_\theta - X$. Thus, in designing the perturbations $\{f_\theta, \theta \in \Theta\}$ for variance reduction, one can exploit the correlation between X and $X_\theta - X$, which is much clearer than that between X and S_θ . See Figure 1 for an illustration. (For the Gaussian case, if X_θ is ‘close’ to X , then S_θ is almost aligned with X ; see Proposition 2 below.)

The unit vector $(X_\theta - X)/\|X_\theta - X\|$ is in the limit (as $\theta \rightarrow \theta_0$) the direction at which the trajectory X_θ approaches X . Proposition 1 implies that if we can choose the family of densities f_θ so that the trajectory X_θ approaches X in a ‘nice’ direction, then the score gives good variance reduction.




Figure 1: A interpretation of S_θ as asymptotically the difference $X_\theta - X$.

In general, one would have to perform a search in Θ to determine the θ^* for which the variance reduction is maximised, i.e.,

$$\theta^* = \arg \min_{\theta \in \Theta} \min_{u \in \mathbf{R}} \text{var}(X(W) + uS_\theta(W)). \quad (27)$$

Additionally, the reduction in variance achieved at θ^* will depend on the density of interest f_{θ_0} and the family of perturbations considered, i.e., $f_\theta, \theta \in \Theta \setminus \{\theta_0\}$. The following result establishes that if f_{θ_0} is Gaussian, then $S_\theta(W)$ is a perfect control variate in the limit as $\theta \rightarrow \theta_0$.

Proposition 2 *For each $\theta \in (0, \infty)$ ($\Theta := (0, \infty)$), let f_θ be the density of a Gaussian random variable with mean μ and variance θ^2 . Let $\theta_0 \in (0, \infty)$ be some fixed parameter. Then, the hypothesis of Proposition 1 is satisfied. Furthermore,*

$$\lim_{\theta \rightarrow \theta_0} \min_{u \in \mathbf{R}} \text{var}(X(W) + uS_\theta(W)) = 0.$$

(All proofs appear in [7].)

4 CV for the KL Criterion

In order to implement S_θ as a CV when estimating the entropy of an arbitrary density f_{θ_0} , we need to first define $\{f_\theta, \theta \in \Theta \setminus \{\theta_0\}\}$. In the Gaussian case, it was shown that if $\{f_\theta, \theta \in \Theta \setminus \{\theta_0\}\}$ were defined to be perturbations of f_{θ_0} in variance only and not the mean, S_θ is a perfect control variate in the limit as $\theta \rightarrow \theta_0$. In general, one would have to perform a search in Θ to determine the θ^* for which the variance reduction is maximised. Furthermore, $\{f_\theta, \theta \in \Theta \setminus \{\theta_0\}\}$ need not be restricted to a perturbation of f_{θ_0} in variance only. This section addresses the problem of implementing these perturbations for the KL criterion in (10).

Defining $\{f_\theta, \theta \in \Theta \setminus \{\theta_0\}\}$ Write (10) as the difference

$$\begin{aligned} & \int \log(q_k(y|x, a)) q_k(y|x, a) \Psi_k[\pi_{k-1}](x) dx dy \\ & - \int \log(\langle q_k(y|\cdot, a), \Psi_k[\pi_{k-1}] \rangle) \\ & \times \langle q_k(y|\cdot, a), \Psi_k[\pi_{k-1}] \rangle dy. \end{aligned} \quad (28)$$

Recall that we have $\{(y^{(i)}, x^{(i)}, L^{-1})\}_{i=1}^L$ (with equal weights) as a particle approximation to $q(y|\cdot)\Psi[\pi]$. (Since we are concerned with a fixed action a and epoch k , they may be omitted from the notation without loss of generality.)

Consider the first term in (28)

$$\int \log(q(y|x)\Psi[\pi](x)) q(y|x)\Psi[\pi](x) dx dy.$$

(The treatment of the second term is similar and is omitted.) Consider first the specific case when the measurement model is the bearings only model

$$y = \arctan\left(\frac{x_1 - p_{1,a}}{x_3 - p_{2,a}}\right) + \eta, \quad (29)$$

where $x = [x_1, x_2, x_3, x_4]^T$ with $[x_1, x_3]^T$ and $[x_2, x_4]^T$ being the position and velocity of the target respectively, $p_a = [p_{1,a}, p_{2,a}]^T$ is the x-y position of the sensor and the noise η is $N(0, \sigma_\eta^2)$. Let $\Theta \subset \mathbf{R}$ contain σ_η . For each $\theta \in \Theta$, define

$$\begin{aligned} & f_\theta(x, y) \\ & = \frac{1}{\theta\sqrt{2\pi}} \exp\left(-\frac{1}{2\theta^2} \left(y - \arctan\left(\frac{x_1 - p_{1,a}}{x_3 - p_{2,a}}\right)\right)^2\right) \\ & \times \Psi[\pi](x) \end{aligned} \quad (30)$$

Thus $f_{\sigma_\eta}(x, y) = q(y|x)\Psi[\pi](x)$. Although the exact density $\Psi[\pi](x)$ is unavailable, the score is well defined and available to be implemented as a CV, i.e.,

$$\begin{aligned} & S_\theta(x, y) + 1 \\ & = \frac{\sigma_\eta \exp\left(-\frac{1}{2\theta^2} \left(y - \arctan\left(\frac{x_1 - p_{1,a}}{x_3 - p_{2,a}}\right)\right)^2\right)}{\theta \exp\left(-\frac{1}{2\sigma_\eta^2} \left(y - \arctan\left(\frac{x_1 - p_{1,a}}{x_3 - p_{2,a}}\right)\right)^2\right)} \end{aligned} \quad (31)$$

This scheme is implemented in the numerical example in Section 5. Note that $f_{\sigma_\eta}(x, y) = \lim_{\theta \rightarrow \sigma_\eta} f_\theta(x, y)$.

The above example illustrates how one may define perturbations $\{f_\theta, \theta \in \Theta\}$ of f_{θ_0} assuming the analytic form of $q(y|x)$ is available. If this were not the case, the following scheme may be used instead.

Let f be the density for which we wish to estimate the entropy. We assume that an analytic expression

for f is unavailable but it can be evaluated or estimated pointwise. Let $V \sim N(0, 1)$ and $W \sim f$. For each $\theta \in \mathbf{R}$, the random variable $W + \theta V \sim \int \frac{1}{\theta} f(x - v) p(\frac{v}{\theta}) dv$ where p is the density of $N(0, 1)$. The mean of $W + \theta V$ is unchanged but its variance has, and we may use the density of $W + \theta V$ as f_θ . Although, we cannot evaluate this convolution exactly, we can approximate it. Let $\{v^{(i)}\}_{i=1}^M$ be samples from $N(0, 1)$. Then,

$$f_\theta(x) = M^{-1} \sum_{i=1}^M f(x - \theta v^{(i)}). \quad (32)$$

Because this is a finite sample approximation to the integral $\int \frac{1}{\theta} f(x - v) p(\frac{v}{\theta}) dv$, its mean does not necessarily coincide with that of f unless $M^{-1} \sum_{i=1}^M v^{(i)} = 0$. Note that $f(x) = \lim_{\theta \rightarrow 0} f_\theta(x)$.

Searching For θ^* Consider the problem of identifying the parameter θ that minimises $\min_{u \in \mathbf{R}} \text{var}(\log f_{\theta_0}(W) + u S_\theta(W))$, i.e., θ^* in (27). From (24), θ^* also satisfies

$$\frac{\langle X, S_{\theta^*} \rangle^2}{\|S_{\theta^*}\|^2} \geq \frac{\langle X, S_\theta \rangle^2}{\|S_\theta\|^2}, \quad \forall \theta \in \Theta \setminus \{\theta_0\}. \quad (33)$$

(Note that $X(W) = \log f_{\theta_0}(W)$ (26).) Let $\{w^{(i)}\}_{i=1}^M$ be samples from $f_{\theta_0}(W)$. A Monte Carlo implementation of the above maximisation is

$$\max_{\theta \in \Theta \setminus \{\theta_0\}} \frac{\left(L^{-1} \sum_{i=1}^L X(w^{(i)}) S_\theta(w^{(i)}) \right)^2}{L^{-1} \sum_{i=1}^L S_\theta(w^{(i)})^2}. \quad (34)$$

This is a continuous optimisation problem when the score is a continuous function of θ . This is certainly the case for (31), and (32) if f is continuous. In implementation, one may simply search a finite subset of Θ for a good candidate. Note that same set of samples $\{w^{(i)}\}_{i=1}^M$ are used when evaluating the objective on Θ . This is important when the source of the samples is a particle filter as in the case of estimating the KL criterion.

5 Numerical Example

Consider a bearings only tracking application where the target of interest moves in the region $[-100, 100] \times [-100, 100]$; see Figure 2. There are 4 available sensors located at $p_1 = [-60, 20]^T$, $p_2 = [0, -40]^T$, $p_3 = [40, -90]^T$, $p_4 = [50, -10]^T$. The state of the target consists of position and velocity, while only its bearings are observed. The states follows a linear dynamics,

$$x_k = \begin{bmatrix} 1 & T & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & T \\ 0 & 0 & 0 & 1 \end{bmatrix} x_{k-1} + \begin{bmatrix} \frac{T^2}{2} & 0 \\ T & 0 \\ 0 & \frac{T^2}{2} \\ 0 & T \end{bmatrix} \begin{bmatrix} v_{1,k} \\ v_{2,k} \end{bmatrix}$$

where $x_k = [x_{1,k}, x_{2,k}, x_{3,k}, x_{4,k}]^T$; $[x_{1,k}, x_{3,k}]^T$ is the position while $[x_{2,k}, x_{4,k}]^T$ is the velocity. The initial state of the target is $x_0 = [-80, 2, 50, 2.7]^T$. The measurement θ_k received by sensor s located at p_s at time k is given by

$$y_k = \begin{cases} \arctan\left(\frac{x_{1,k} - p_{1,s}}{x_{3,k} - p_{2,s}}\right) + \eta_k, & x_{3,k} - p_{2,s} > 0 \\ \zeta_k \sim [0, \pi], & \text{otherwise} \end{cases}$$

The noise processes $\{v_{1,k}\}$, $\{v_{2,k}\}$, and $\{\eta_k\}$ are mutually independent zero-mean Gaussian white noise with respective standard deviations $\sigma_{v_1} = \sigma_{v_2} = 1$, $\sigma_\eta = 0.01$ (i.e. approximately 0.6 of a degree). Note that when the target is outside the field of view (FoV), i.e. when $x_{3,k} - p_{2,s} \leq 0$, the measurement θ_k is uniformly distributed between $[0, \pi]$. Hence, sensors can only look 'upwards'.

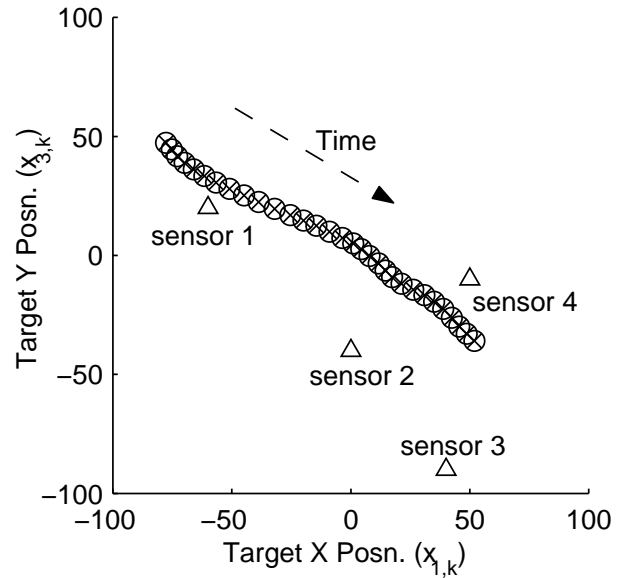


Figure 2: True target track and sensor positions.

Figure 2 shows the true track positions and the sensor locations. Shown in Figure 3 is the reduction in the variance of the KL objective estimate. All the plots were generated with 250 particles. The score defined in (31) was implemented as the CV where the set Θ comprised of 36 elements; each $\theta \in \Theta$ was a small perturbation of σ_η . A search over Θ was performed until an element θ was found for which the variance reduction exceeded 80%; the optimum θ^* as defined in (34) was not sought. Figure 3 shows the variance reduction for both the numerator and the denominator of the KL objective separately for each of the 4 sensors. Note that in most cases the CV method reduced the variance by about 80%; the worse case reduction was 40% and thus more than doubling the effective number of particles. Figure 4 depicts the filtering density

of the targets position at different times when sensor scheduling is applied. The active sensor is overlaid with an '*'.

Note that searching over Θ for θ^* is simple as (34) is a closed-form expression in θ . Naturally, it would also be possible to achieve variance reduction by simply increasing the number of particles. Our premise is that this option is not available or, for a given set of particles, the aim is to extract the most accurate estimates of the quantities of interest. The proposed CV method is non-invasive unlike an importance sampling based variance reduction method, as it does not require any modification to the process of generating the samples (or particle filter). Rather, it only collects additional statistics from these samples and it is very simple to implement. Furthermore, we have demonstrated in Figure 3 that at some instances, the proposed method reduces the variance to less than 10% of the original – to achieve $x\%$ of the original variance, requires $100/x$ fold increase in the number of particles.

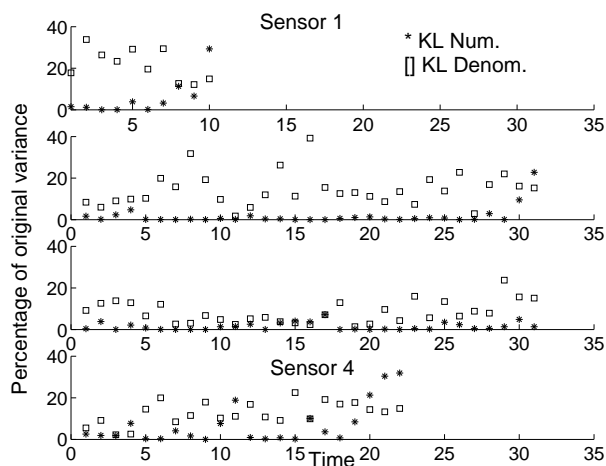


Figure 3: Variance reduction achieved for all sensors. KL numerator and denominator displayed separately. Truncated plots correspond to the target being out of the FoV of the sensor.

6 Conclusion

We have presented a Control Variate scheme for variance reduction in SMC based one-step ahead sensor management. It was shown for the Gaussian case that one can construct an optimal CV asymptotically and in the numerical example, large variance reduction was demonstrated. However, these results are still preliminary as many questions remain unanswered. For example, it would be of interest to similarly quantify the variance reduction for other analytic densities and the more general Gaussian mixture, as the latter is frequently used in target tracking. The proposed CV

scheme may also be used when we have a horizon- N formulation of the sensor management problem [6], where one implements open (and closed)-loop feedback control policies with SMC. This avenue is currently being investigated. We are also studying how one should define the perturbations $\{f_\theta, \theta \in \Theta\}$ in order to achieve large variance reduction.

References

- [1] D.P. Bertsekas, *Dynamic programming and optimal control*. Belmont: Athena Scientific, 1995.
- [2] P.W. Glynn and R. Szechtman, “Some new perspectives on the method of control variates,” in *Monte Carlo and Quasi-Monte Carlo Methods 2000*, K.T. Fang, F.J. Hickernell and H. Niederreiter, Eds., Springer-Verlag, Berlin, pp. 27–49, 2002.
- [3] S.G. Henderson and B. Simon, “Adaptive simulation using perfect control variates,” 2003. Available: <http://www.orie.cornell.edu/~shane/>
- [4] A. Doucet, B. Vo, C. Andrieu and M. Davy, “Particle filters for multi-target tracking and sensor management,” in *Proc. 5th Annual Conference on Information Fusion*, vol. 1, pp. 474–481, Annapolis, Maryland 2002.
- [5] A. Doucet, J.F.G. de Freitas and N.J. Gordon (Eds.), *Sequential Monte Carlo methods in practice*. New York: Springer, 2001.
- [6] S. Singh, B. Vo, A. Doucet and R.J. Evans, “Stochastic approximation for optimal observer trajectory planning in bearings-only tracking,” *IEEE Conf. Decision Control*, 2003.
- [7] S. Singh, B. Vo, A. Doucet and R.J. Evans, “Sequential Monte Carlo for finite horizon general state space POMDPs,” in preparation, 2004.




Figure 4: Tracking with sensor scheduling: the peaks of the filtering densities are very close to the true target. Active sensor overlaid with '*'. (If document viewed as a soft-copy, zoom in for a closer view of the particles.)