# Sequential Importance Sampling for Visual Tracking Reconsidered


**Péter Torma**
Mindmaker, Ltd.
Budapest 1121 HU
Konkoly Th. M. u. 29-33
tyus@mindmaker.hu

**Csaba Szepesvári**
Mindmaker, Ltd.
Budapest 1121 HU
Konkoly Th. M. u. 29-33
szepes@mindmaker.hu



## Abstract

We consider sequential importance sampling for filtering dynamical systems observed in noise and when the importance function is defined over a few selected components of the state space, typically the components corresponding to the innovation part of the process to be filtered. In this case the basic importance sampling algorithm yields high variance estimates of the posterior. The problem appears for example in visual tracking using particle filters when one uses an importance function based on the output of color blob detector. In this paper we propose several new methods that are proven to be unbiased and show their increased efficiency.


## 1 Introduction

Let us consider sequential importance sampling algorithms for the estimation of the posterior $p(x_t|y_{0:t})$ corresponding to the observation sequence $y_{0:t} = (y_0, \ldots, y_t)$ ($y_t \in \mathbf{R}^m$) of the dynamical system with state $x_t$ evolving in time ($x_t \in \mathbf{R}^n$) according to a generic Markov model defined by the transition kernel $p(x_t|x_{t-1})$ and where the observation density is given by $p(y_t|x_t)$.

In this paper we shall assume the followings:

**Assumption A** Partitioned dynamical model. Assume that the state $x_t$ is partitioned into two parts:

$$x_t = \begin{pmatrix} x_{1,t} \\ x_{2,t} \end{pmatrix} \quad (1)$$

such that its evolution is given by

$$x_{1,t+1} = f_1(x_t) + w_t, \quad (2)$$
$$x_{2,t+1} = f_2(x_t), \quad (3)$$

i.e., the component $x_{1,t}$ evolves stochastically, while $x_{2,t}$ evolves deterministically. For obvious reasons, $x_{2,t}$ will be called the history part and $x_{1,t}$ will be called the innovation part of the state.

**Assumption B** Proposal density is defined over the stochastic component. In more details, assume that the proposal density $\pi$ used in the importance sampling step is defined over the subspace of the state space corresponding to the first component $x_{1,t}$ of the state.

**Assumption C** Cost of computing the evaluation of the observation density is high.

In order to further simplify the exposition we shall assume the following:

**Assumption D** The observation density depends only on $x_{1,t}$, the innovation part of the state:

$$p(y_{t+1}|x_{t+1}) = p(y_{t+1}|x_{1,t+1})$$

(where the notation is abused as is usual in the literature).

Assumptions A, B, and C are quite often satisfied in visual tracking problems. First, AR processes are often used in visual tracking problems and AR processes satisfy Assumption A since their dynamics is given by

$$\begin{aligned} x_{1,t+1} &= \textstyle\sum_{i=0}^{p-1} A_i \cdot (x_t)_{i+1} + Be_{t+1}, \\ (x_{2,t+1})_j &= (x_t)_j, \quad j = 1, 2, \ldots, p-1. \end{aligned}$$

Here $(x_t)_i$ denotes the $i$th block of $x_t$. It is also quite common that Assumption B is satisfied e.g. when color blob detection is used to produce the proposal density as in [3]. Finally, Assumption C is justified since in particle filter based visual tracking it is the evaluation of the observation density that involves the expensive image processing steps.

Sequential importance sampling has been considered in the literature by many authors (for an overview see [2]). The main loop of the classical particle filter algorithm adapted to our assumption is shown in

Table 1. Here, for the sake of simplicity, we have presented the variant that resamples the particle set in each step and we have assumed that the proposal density depends only on the most recent observation and not on the past states – this is also a relevant assumption in visual tracking when the proposal density depends only the last frame observed. In the followings the initialization phase of the algorithms are omitted as they are usually trivial. Provided

1. Sample $u_{t+1}^{(i)} \sim \pi(u|y_{t+1})$

   and let $\hat{x}_{t+1}^{(i)} = \begin{pmatrix} u_{t+1}^{(i)} \\ f_2(x_t^{(i)}) \end{pmatrix}$.

2. Let $w_{t+1}^{(i)} \propto w_t^{(i)} \frac{p(y_{t+1}|\hat{x}_{t+1}^{(i)})p(\hat{x}_{t+1}^{(i)}|x_t^{(i)})}{\pi(u_{t+1}^{(i)}|y_{t+1})}$).

3. Sample $j_{t+1}^{(i)} \sim (w_{t+1}^{(1)}, \ldots, w_{t+1}^{(N)})$.

4. Let $x_{t+1}^{(i)} = \hat{x}_{t+1}^{(j_{t+1}^{(i)})}$.

Table 1: Main Loop of Basic SIR. See e.g. [2]. $N$ is the number of particles and $i = 1, 2, \ldots, N$ is a particle index.

that the weights and particles are initialized appropriately, the particle set $x_t^{(i)}$ can be shown to represent an unbiased estimate of the posterior, i.e., for any measurable function $h$ defined over the state space, $E\left[(1/N)\sum_{i=1}^N h(x_t^{(i)}) \,|\, y_{0:t}\right] = E\left[h(x_t) \,|\, y_{0:t}\right]$. Unfortunately, however, this algorithm can be very inefficient, and may require a large number of particles to achieve even a modest precision, as in Step 2 many weights can get small quickly. This is because the multiplicative term $p(\hat{x}_{t+1}^{(i)}|x_t^{(i)})$ can be very small when $p(u_{t+1}^{(i)}|x_t^{(i)})$ is small. Since $u_{t+1}^{(i)}$ is sampled independently of the past of the process and so it can be "contradicting" with the history of the process this can be expected to be the case. This is illustrated below on Figure 1.

The inefficiency of particle filter methods have been observed by many authors. The typical solution is to choose a better proposal distribution. For example, Pitt and Shephard proposed the so-called auxiliary variable method whereas the proposal distribution is an appropriately defined mixture that depends both on the past state and the most recent observation [5]. Although their method overcomes the problem of weight degeneracy in most of the cases, it involves $R \gg N$ evaluation steps of the observation density and thus the computational burden of this algorithm when applied to visual tracking can be pretty high.

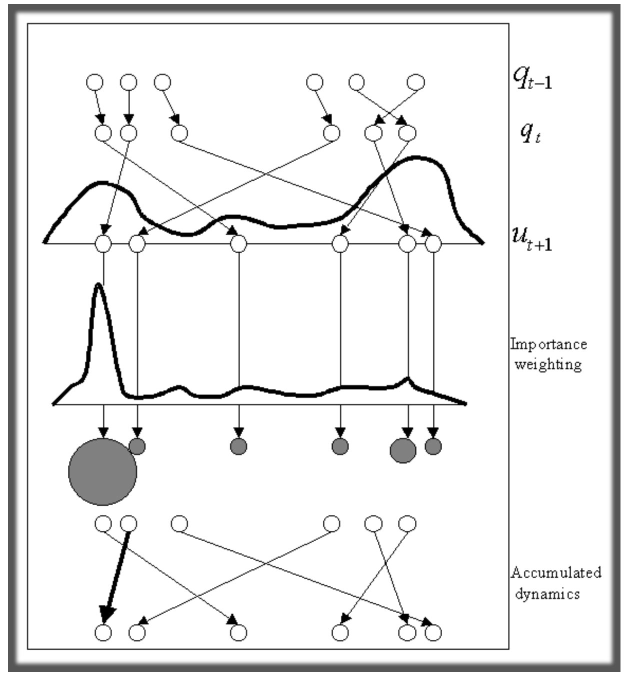


Figure 1: Illustration of the behavior of the Basic SIR Algorithm. Assume that the state is two dimensional and evolves according to a linear model. The first two lines show the state $x_t = (q_t, q_{t-1})$. Next, the proposal density is shown and the innovations drawn from it ($u_{t+1}$). The arrows from $q_t$ to $u_{t+1}$ show the association of innovations with their histories. The importance weights will prefer likely associations, but the number of these can be very low. In the final part, the importance of a pair is shown by the strength of the arrow pointing from the past configuration towards the associated innovation component.

More recently, van der Merwe et. al proposed to use a bank of unscented Kalman filters to define the proposal distribution [8]. This algorithm, called the unscented particle filter, is similar in essence to the auxiliary variable method but avoid the expensive likelihood calculations. Yet another method is presented in [7], where the authors propose to trade off bias against variance by adjusting the location of particles after the prediction step so that the observation density is locally maximized by the particles after the adjustment.

Here we shall attempt to propose new efficient algorithms for the case where the proposal is restricted to have the form $\pi(x_{1,t}|y_{t+1})$. Therefore we have:

**Assumption E** The proposal density, which is in general $\pi(x_{t+1}|y_{t+1}, x_{0:t})$, takes the form $\pi(x_{1,t}|y_{t+1})$.

This is an attractive assumption since importance functions of the above form are easy to define and sample from. Further, this is exactly the case that was considered in the computational example on which Isard and Blake illustrated their ICondensation algorithm [3]. In this sense ICondensation is the algorithm that is the most directly relevant to our case. We will discuss ICondensation later after introducing the new algorithms.

## 2 Algorithms

The idea of the algorithms we consider is to ensure that for each particle the history part of the particle will match the innovation component sampled from the proposal. Effectively, we solve this by drawing an appropriate history for each innovation sampled.

The main loop of the first algorithm we propose is given in Table 2.

1. Sample $u_{t+1}^{(i)} \sim \pi(u_{t+1}|y_{t+1})$
2. Sample $J_{t+1}^{(i)} \sim$
$$\left(\ldots, \frac{p(y_{t+1}|u_{t+1}^{(\cdot)})}{\pi(u_{t+1}^{(\cdot)}|y_{t+1})} \sum_{k=1}^{N} w_t^{(k)} p(u_{t+1}^{(\cdot)}|x_t^{(k)}), \ldots\right)$$
3. Sample $I_{t+1}^{(i)} \sim (\ldots, w_t^{(\cdot)} p(u_{t+1}^{(J_{t+1}^{(i)})}|x_t^{(\cdot)}), \ldots)$.
4. Let $x_{t+1}^{(i)} = \begin{pmatrix} u_{t+1}^{(J_{t+1}^{(i)})} \\ f_2\left(x_t^{(I_{t+1}^{(i)})}\right) \end{pmatrix}$.

Table 2: Main Loop of SIR with History Sampling.

In order to understand this algorithm, let us introduce the auxiliary variables $(x_t^{(i,j)}, w_t^{(i,j)})$ such that $(x_t^{(i,j)}, w_t^{(i,j)}) = (x_t^{(i)}, w_t^{(i)})$. Let the particle set at time $t+1$ be defined by the equations

$$\begin{aligned} x_{t+1}^{(i,j)} &= \begin{pmatrix} u_{t+1}^{(j)} \\ f_2(x_t^{(i)}) \end{pmatrix} \\ w_{t+1}^{(i,j)} &= w_t^{(i,j)} \frac{p(y_{t+1}|x_{t+1}^{(i,j)}) p(x_{t+1}^{(i,j)}|x_t^{(i,j)})}{\pi(u_{t+1}^{(j)}|y_{t+1})} \\ &= w_t^{(i)} \frac{p(y_{t+1}|u_{t+1}^{(j)}) p(u_{t+1}^{(j)}|x_t^{(i)})}{\pi(u_{t+1}^{(j)}|y_{t+1})}, \end{aligned}$$

where the last equality follows by our assumptions on the observation and proposal densities. Now assume that at time $t$ the particle set $(x_t^{(i)}, w_t^{(i)})_{i=1}^N$ represents an unbiased estimate of the posterior $p(x_t|y_{0:t})$. Clearly, by the unbiasedness of the basic importance sampling scheme, the particle set $(x_{t+1}^{(i,j)}, w_{t+1}^{(i,j)})_{i,j=1}^N$ will represent an unbiased estimate of the posterior $p(x_{t+1}|y_{0:t+1})$ (cf. Steps 1–2 of the basic SIR algorithm of Table 1). Now, if $I_{t+1}^{(i)}$ and $J_{t+1}^{(i)}$ are the random indexes drawn as in Step 2 and 3 of the algorithm, then

$$P(I_{t+1}^{(i)} = k, J_{t+1}^{(i)} = l) = P(I_{t+1}^{(i)} = k | J_{t+1}^{(i)} = l)$$
$$P(J_{t+1}^{(i)} = l) = \frac{w_t^{(k,l)} p(u_{t+1}^{(l)}|x_t^{(k)})}{\sum_{l=1}^N w_t^{(k,l)} p(u_{t+1}^{(l)}|x_t^{(k)})}$$
$$\frac{\sum_{l=1}^N w_t^{(k,l)} p(u_{t+1}^{(l)}|x_t^{(k)})}{\sum_{k=1}^N \sum_{l=1}^N w_t^{(k,l)} p(u_{t+1}^{(l)}|x_t^{(k)})} \propto w_{t+1}^{(k,l)}$$

Therefore Steps 2 and 3 of algorithm take the form of a standard resampling step and therefore the particle set $(x_{t+1}^{(i)}, 1/N)_{i=1}^N$ as defined above will also give an unbiased estimate of the posterior. Actually, Step 2 and 3 of the above algorithm can be considered as sampling from $(\ldots, w_{t+1}^{(\cdot,\cdot)}, \ldots)$ by means of partitioned sampling [4]: Step 2 samples the innovations to use, whilst Step 3 samples the appropriate histories to be associated with these innovation components.

The advantage of the algorithm against the basic SIR algorithm should be clear: the algorithm selects (by random sampling) pairs of innovations and histories that have high probability of co-occurring, by reducing the variance of estimates. Convergence theorems similar to those of [1] can be derived but are omitted due to lack of space.

The next algorithm we consider is a variant where the innovations are not sampled. This algorithm can be considered as a Rao-Blackwellised version of the previous one whereas the unnecessary steps of sampling the innovation is avoided, thereby further reducing the variance. The algorithm is shown in Table 3.

1. Sample $u_{t+1}^{(i)} \sim \pi(u_{t+1}|y_{t+1})$
2. Sample $I_{t+1}^{(i)} \sim (\ldots, w_t^{(\cdot)} p(u_{t+1}^{(i)}|x_t^{(\cdot)}), \ldots)$.
3. Let $x_{t+1}^{(i)} = \begin{pmatrix} u_{t+1}^{(i)} \\ f_2\left(x_t^{(I_{t+1}^{(i)})}\right) \end{pmatrix}$.
4. Compute $w_{t+1}^{(i)} =$
$$\frac{p(y_{t+1}|u_{t+1}^{(i)}) \sum_{l=1}^N w_t^{(l)} p(u_{t+1}^{(i)}|x_t^{(l)})}{\pi(u_{t+1}^{(i)}|y_{t+1})}.$$

Table 3: Main Loop of RB-SIR with History Sampling.

This algorithm is illustrated on Figure 2. Clearly, in

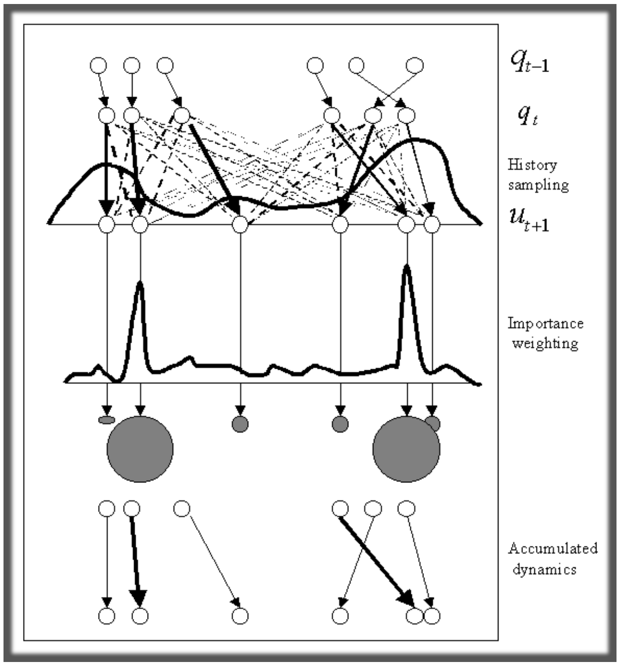


Figure 2: Illustration of the behavior of the RB-SIR History Sampling Algorithm. Again, we assume that the state is two dimensional and evolves according to a linear model. The first two lines show the state $x_t = (q_t, q_{t-1})$. Next the proposal density is shown and the innovations $u_{t+1}$ drawn from it. The arrows from $q_t$ to $u_{t+1}$ show the association of innovations with their histories. The strength of these arrows are proportional to the weights that are used in sampling the history for each innovation. Finally, the importance weighting function is shown. The final set of particles is shown in the last two lines, where again the strength of the arrows is proportional to the weight of the given particle.

contrast to the basic SIR algorithm, the effective sample size will stay higher as sampling the histories helps to reduce the effect of arbitrary innovation-history associations.

In order to show the unbiasedness of the algorithm, in this case we evaluate $R = E[\sum_{i=1}^N w_{t+1}^{(i)} h(x_{t+1}^{(i)}) | y_{0:t+1}]$ directly:

$$
\begin{aligned}
R &= E[\sum_{i=1}^N w_{t+1}^{(i)} h(x_{t+1}^{(I_{t+1}^{(i)},i)}) \mid y_{0:t+1}] \\
&= \sum_{i=1,l=1}^N E\Big[ P(I_{t+1}^{(i)} = l | y_{0:t+1}, w_t^{(\cdot)}, x_t^{(\cdot)}, u_{t+1}^{(\cdot)}) \cdot \\
&\quad E\big[ w_{t+1}^{(i)} h(x_{t+1}^{(I_{t+1}^{(i)},i)}) \mid y_{0:t+1}, \\
&\quad I_{t+1}^{(i)} = l, w_t^{(\cdot)}, x_t^{(\cdot)}, u_{t+1}^{(\cdot)}\big] \mid y_{0:t+1} \Big]
\end{aligned}
$$

Now, since

$$
\begin{aligned}
&P(I_{t+1}^{(i)} = l | y_{0:t+1}, w_t^{(\cdot)}, x_t^{(\cdot)}, u_{t+1}^{(\cdot)}) = \\
&= \frac{w_t^{(l)} p(u_{t+1}^{(i)} | x_t^{(l)})}{\sum_{r=1}^N w_t^{(r)} p(u_{t+1}^{(i)} | x_t^{(r)})}
\end{aligned}
$$

and plugging in the definition of $w_{t+1}^{(i)}$, we get

$$
\begin{aligned}
R = & \sum_{l=1,i=1}^N E\Bigg[ \frac{w_t^{(l)} p(u_{t+1}^{(i)} | x_t^{(l)})}{\sum_{r=1}^N w_t^{(r)} p(u_{t+1}^{(i)} | x_t^{(r)})} \cdot \\
& \frac{p(y_{t+1} | u_{t+1}^{(i)}) \sum_{k=1}^N w_t^{(k)} p(u_{t+1}^{(i)} | x_t^{(k)})}{\pi(u_{t+1}^{(i)} | y_{t+1})} \cdot \\
& h(x_{t+1}^{(l,i)}) \mid y_{0:t+1} \Bigg] = \sum_{l=1,i=1}^N E\Bigg[ w_t^{(l)} p(u_{t+1}^{(i)} | x_t^{(l)}) \cdot \\
& \frac{p(y_{t+1} | u_{t+1}^{(i)})}{\pi(u_{t+1}^{(i)} | y_{t+1})} \cdot h(x_{t+1}^{(l,i)}) \mid y_{0:t+1} \Bigg] \\
& = \sum_{l=1,i=1}^N E\Big[ w_{t+1}^{(l,i)} h(x_{t+1}^{(l,i)}) \mid y_{0:t+1} \Big],
\end{aligned}
$$

which finishes the proof of the claim, since the particle set $(x_{t+1}^{(p,q)}, w_{t+1}^{(p,q)})$ is known to give an unbiased representation of the posterior.

Again, under appropriate conditions, bounds on the variance of the estimate of the posterior can be derived along the lines of [1].

Other variants of the algorithm can be given. As an example let us mention the variant when the history component is kept for each particle and the innovation component is resampled. This variant would only be useful if one believes that the particle set bears more information about the posterior than the proposal function. Since it is quite typical that one believes in the opposite of this, we do not consider this variant here.

Finally, let us mention the practical case when the proposal function is further restricted to selected components of the innovation. For example in visual tracking, often the proposal function is defined only for the translational components and not for the other deformational parameters (e.g. rotation). This is exactly the case of the blob detection defined by Isard and Blake [3]. In this case one can still use a variant of RB-SIR with History Sampling. The main loop of this algorithm is given in Table 4.

We shall still use $u_{t+1}$ for the component sampled from the proposal, while $v_{t+1}$ is the subspace of $x_{1,t+1}$ which

| | |
|---|---|
| 1. | Sample $u_{t+1}^{(i)} \sim \pi(u_{t+1}\|y_{t+1})$ |
| 2. | Sample $I_{t+1}^{(i)} \sim (\ldots, w_t^{(\cdot)} p(u_{t+1}^{(i)}\|x_t^{(\cdot)}), \ldots)$. |
| 3. | Draw $v_{t+1}^{(i)}$ from $p(v_{t+1}^{(i)}\|u_{t+1}^{(i)}, x_t^{(I_{t+1}^{(i)})})$. |
| 4. | Let $x_{t+1}^{(i)} = \begin{pmatrix} u_{t+1}^{(i)} \\ v_{t+1}^{(i)} \\ f_2\left(x_t^{(I_{t+1}^{(i)})}\right) \end{pmatrix}$. |
| 4. | Compute $w_{t+1}^{(i)} = \frac{p(y_{t+1}\|u_{t+1}^{(i)}) \sum_{l=1}^N w_t^{(l)} p(u_{t+1}^{(i)}\|x_t^{(l)})}{\pi(u_{t+1}^{(i)}\|y_{t+1})}$. |

Table 4: Main Loop of RB-subspace-SIR with History Sampling.

cannot be determined by the importance function.[1] For the derivation on the importance weight note that:

$$p(x_{1,t+1}^{(i)}|x_t^{(I_{t+1}^{(i)})}) = p(u_{t+1}^{(i)}, v_{t+1}^{(i)}|x_t^{(I_{t+1}^{(i)})}) =$$
$$= p(v_{t+1}^{(i)}|u_{t+1}^{(i)}, x_t^{(I_{t+1}^{(i)})}) p(u_{t+1}^{(i)}|x_t^{(I_{t+1}^{(i)})}).$$

Here the first term is the one $v_{t+1}^{(i)}$ is drawn from in Step 3 of the algorithm, so only the second term remains in the importance weight.

Sampling from $p(v_{t+1}|u_{t+1}, x_t)$ is not always straightforward, but it is in the most usual practical cases. For example if the system noise is Gaussian, it is a marginal of it, still being Gaussian. In many other cases $p(v_{t+1}|u_{t+1}, x_t) = p(v_t|x_t)$, which certainly further simplifies the algorithm.

## 3 Experiments

To purpose of this section is to present some experimental results proving the effectiveness of the new algorithms. The algorithms were tested on a contour based visual tracking problem where the dynamical model was a second-order AR process and the observation density is defined using local measurements on the image at the normals of the spline representing the contour to be tracked. The four dimensional configuration space is the Descartes product of the spaces of the translation, rotation and scale parameters. Details of the computations can be found in [6].

[1] In Blake and Isard [3] $u_{t+1}$ is the translation subspace, while $v_{t+1}$ is called deformation subspace.

For the proposal we used a Gaussian color blob detector working on the original frames (the resolution of the original frames are $240 \times 180$) and then the output was down-sampled to a resolution of $24 \times 18$. Spatial indexes were then drawn from the appropriately rescaled output of the blob detection and the final random spatial positions are obtained by applying a random perturbation to the corresponding central on the image positions. This was done by sampling another "fine-scale" index uniformly over a $10 \times 10$ matrix and using the obtained values to "refine" the initiate positions. The task was to track an artificial object, while another object with the same color and shape was lying on the table to make the task of tracking non-trivial; i.e. "tracking in clutter". For the sake of comparisons CONDENSATION (also known as N-IPS), the basic SIR algorithm and RB-subspace-SIR with History Sampling (the third variant) were tried. A typical tracking scenario using RB-subspace-SIR with History Sampling is presented in Figure 3.

In one sequence of 5 seconds of video sampled at 30 fps the configurations of the object to be tracked were determined manually (this is the sequence shown on Figure 3). Each algorithm was then tested on this sequence with 100 different random initializations. Accuracy of tracking and the probability of losing the object to be tracked were measured. Equivalent running time experiments were considered on an Intel Pentium IV 1.4GHz computer with 128M RAM, i.e., the particle sizes were set so that the running time of the algorithms were the same (no special attempt was made to optimize any of the algorithms). The corresponding particle sizes are given in Table 5.

| **Algorithm** | **Particle Size** |
|---|---|
| CONDENSATION | 2000 |
| Basic SIR | 3000 |
| RB-subspace-SIR with HS | 400 |

Table 5: Particle sizes used in the experiments.

Results of accuracy measurements are shown in Figure 4. The curves represent the accuracy as the distance to the object to be tracked in pixels averaged over the 100 runs. It should be clear from the figure that in terms of their accuracy the algorithms are ordered as CONDENSATION $\approx$ RB-subspace-SIR with HS $\ll$ Basic SIR.

The probabilities of losing the object are shown in Figure 5. The probability of losing the object is measured by computing the fraction of cases when the hypothesis of the tracker is outside of a certain large neighborhood of the object to be tracked (here defined to be 50 pixels).[2] Again, the ordering of the algorithms remains the same: RB-subspace-SIR with History Sampling tracks the object reliable just as CONDENSATION, while Basic SIR performs much weaker.

By a more close inspection of the data we observed that when CONDENSATION is locked on the object, it quite accurate while if the object is lost the recovery is very slow. On the other hand RB-subspace-SIR with History Sampling recovers fast from temporal mistakes, which is clearly due to the importance sampling scheme.

## 4 Discussion

The above experiments show that the algorithms proposed here to overcome the problems with the efficiency of Basic SIR work as expected. Although the results obtained are very encouraging, one should not forget that the new algorithms require evaluation of dynamical likelihood $N^2$ times (such as in Blake and Isard [3]). Fortunately, the most computationally expensive part, the number of evaluations of the observation densities is still $O(N)$, so in practice, with an efficient implementation[3] the new algorithms can represent a viable alternative to running basic SIR or CONDENSATION with a larger number of particles. This has also been proven here by the experiments.

What remains is the discussion of the relation to ICondensation, the algorithm introduced in [3]. At a first glance ICondensation looks very similar to Basic SIR. However, let us take a closer look at this algorithm. In Figure 1 of [3] in Step 2(a) the next state is sampled from the proposal (as in the Basic SIR algorithm), but in calculating the importance weights Isard and Blake used the calculations given in RB-SIR with History Sampling (Step 2(b) and 3 of their algorithm). The same problem appears when they describe the details of the algorithm in Section 4.2. Clearly ICondensation is not unbiased. Why does this algorithm work at all? One explanation is that when all the particles are concentrated into the same portion of the state space then the importance weights by summing up the prediction densities will be close to the importance weights computed as given in the Basic SIR algorithm. Another explanation is that because of a "mixture" of lucky choice of parameters, models (e.g. The bias introduced will be small if the state prediction density is very flat), a very good proposal density and also because the algorithm mixes CONDENSATION with importance sampling and re-initialization in a clever way. We think that CONDENSATION with re-initialization could achieve similar results. Also, there are results showing that CONDENSATION is robust against perpetual perturbations provided that these perturbations satisfy certain conditions (see Theorem 3.1 of [7]).

[2] In order to separate the effect of losing the object from problems with accuracy when the object is tracked, when the object is lost at a certain point in time, distance measurement is not used in measuring the accuracy.

[3] If the dynamical system noise is Gaussian, one can evaluate the dynamical likelihood with a help of a lookup-table.

## 5 Conclusions

In this paper we have reconsidered sequential importance sampling for filtering dynamical systems when the proposal distribution takes a special form, found commonly in visual tracking problems. We have argued that the "Basic SIR" algorithm can be very inefficient in this case since the probability of the innovation component drawn from the proposal given its associated history can be very low. We have proposed new algorithms with the aim to overcome the problems. Given the innovations sampled from the proposal, the new algorithms sample the "history". The new algorithms were shown to yield unbiased estimates of the posterior and, by means of computer experiments, were shown to yield superior performance as compared with Basic SIR. Further the new algorithm has shown to perform equivalent to CONDENSATION with less particles. Further work shall include a more comprehensive comparison of the new algorithms to other algorithms and in particular a comparison with UPF, and combination of the new algorithms with LS-N-IPS.

## References


[1] D. Crisan and A. Doucet. Convergence of sequential monte carlo methods, 2000.

[2] Arnaud Doucet. On sequential simulation based methods for Bayesian filtering. *Statistics and Computing*, 10(3):197–208, 1998.

[3] Michael Isard and Andrew Blake. ICondensation: Unifying low-level and high-level tracking in a stochastic framework. *Proc 5th European Conf. Computer Vision*, 1998.

[4] John MacCormick. *Probabilistic Modelling and Stochastic Algorithms for Visual Localisation and Tracking.* PhD thesis, University of Oxford, 2000.

[5] Michael K. Pitt and Neil Shephard. Filtering via simulation: Auxiliary particle filter. *Journal of the American Statistical Association*, 94:590–9, 1999.

[6] Csaba Szepesvári and Péter Torma. Vision based object tracking using LS-N-IPS. *In Proc. ISAP'01*, pages 277–282, 2001.

[7] Péter Torma and Csaba Szepesvári. LS-N-IPS: an improvement of particle filters by means of local search. *In Proc. Non-Linear Control Systems(NOLCOS'01)*, 2001.

[8] Rudolph van der Merwe, Nando de Freitas, Arnaud Doucet, and Eric Wan. The unscented particel filter. In *Advances in Neural Information Processing Systems 13*, Nov 2001.

Figure 3: A typical tracking sequence with a 600-particle RB-subspace-SIR with History Sampling. Black contours show probable configurations, while the white contour represents the expected configuration. The object on the table has the same characteristics (e.g. is made of the same material of the same color) as the object to be tracked.

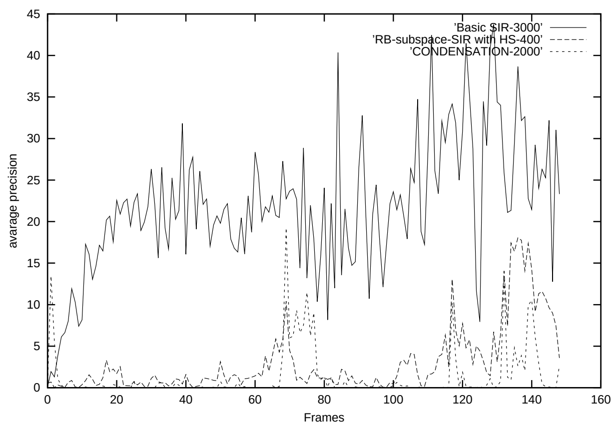


Figure 4: Tracking accuracy as a function of frame number.

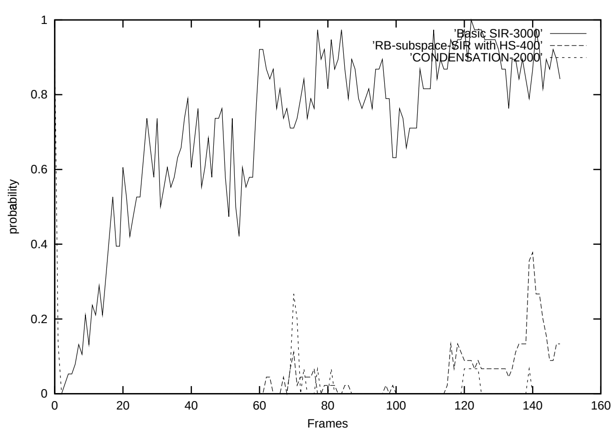


Figure 5: Probability of losing the object.