A Rao-Blackwellized Particle Filter for EigenTracking

Zia Khan, Tucker Balch, and Frank Dellaert
Georgia Institute of Technology College of Computing
{*zkhan,tucker,dellaert*}*@cc.gatech.edu*

Abstract

Subspace representations have been a popular way to model appearance in computer vision. In Jepson and Black's influential paper on EigenTracking, they were successfully applied in tracking. For noisy targets, optimization-based algorithms (including EigenTracking) often fail catastrophically after losing track. Particle filters have recently emerged as a robust method for tracking in the presence of multi-modal distributions. To use subspace representations in a particle filter, the number of samples increases exponentially as the state vector includes the subspace coefficients. We introduce an efficient method for using subspace representations in a particle filter by applying Rao-Blackwellization to integrate out the subspace coefficients in the state vector. Fewer samples are needed since part of the posterior over the state vector is analytically calculated. We use probabilistic principal component analysis to obtain analytically tractable integrals. We show experimental results in a scenario in which we track a target in clutter.

1. Introduction

Subspace representations [21, 2, 6] have been a long-standing and popular way to model appearance and shape in computer vision. These methods model the density over a high-dimensional space of feature vectors using a generative model, where each vector is assumed to be a corrupted version of a linear combination of a small set of basis vectors. The basis vectors are typically obtained by applying principal components analysis (PCA) to a large training set. This approach has been used extensively in a variety of settings, including modeling appearance in face detection and recognition (see e.g. [14] for a review), modeling 2D shape and appearance [7, 6], and 3D shape and appearance [22].

Subspace representations were also successfully used for tracking, e.g. in Jepson and Black's influential paper on EigenTracking [3] in an optimization-based tracking framework. Similarly, Cootes et al [6, 8] model both appearance and shape in a recursive state estimation framework. Unfortunately, for noisy targets, optimization-based tracking algorithms often fail catastrophically after losing track. Hence, *particle filters* have recently emerged as a simple and robust method for tracking in the presence of substantial non-normal measurements and/or dynamics [10, 11, 4]. The particle filter approximates the distribution over the current target state as a set of weighted samples, which is recursively updated based on the current measurement, a target motion model, and a measurement model.

Incorporating subspace representations as the measurement model in a particle filter does present some challenges, however. The simplest approach to incorporating subspace representations into the particle filter framework is to augment the state space of the target with the subspace coefficients. This is the method employed by Blake and Isard to track the an articulated object such as a hand in clutter [12]. However, this proves problematic as the number of samples in the particle filter needs to increase exponentially with the dimensionality of the state space, which now includes the subspace coefficients [13].

In this paper, we introduce an efficient method for using subspace representations as part of a particle filter. In particular, we propose to use Rao-Blackwellization [16, 5] to integrate out the appearance subspace coefficients of the state vector, leaving only the original target state. The advantage of this is that fewer samples are needed since part of the posterior over the state is analytically calculated, rather than being approximated using a more expensive sample representation. We use probabilistic principal component analysis (PPCA), a probabilistic subspace model for which the integral can be computed analytically [20, 17].

Rao-Blackwellized particle filters have been applied in other state estimation problems. In [18], the authors integrate over the 2D target positions and sample over measurement target assignments to track people based on noisy position measurements from IR sensors. In [9], de Freitas uses a Rao-Blackwellized particle filter for fault detection where Kalman filters are applied over continuous parameters and samples obtained over discrete fault states. And finally, in [15], the authors inte-

grate over landmark locations in a robotics application where the goal is to localize a robot while simultaneously building a map of the environment.

In the remainder of this paper we first review the Bayes filter and the particle filter in Section 2, discuss Rao-Blackwellization and Rao-Blackwellized particle filters in Section 3, and then specialize that general framework to Rao-Blackwellized EigenTracking in Section 4. In Section 5 we illustrate our approach in a tracking scenario in which we track a target in clutter. In this scenario, we compare the performance of a standard particle filter to a Rao-Blackwellized particle filter.

2. Bayesian Target Tracking

2.1. The Bayes Filter

Target tracking problems can be expressed as a Bayes filter, in which we recursively update the posterior distribution $P(X_t|Z^t)$ over the target state X_t given all observations $Z^t = \{Z_1..Z_t\}$ up to and including time t, according to:

$$P(X_t|Z^t) \propto P(Z_t|X_t) \int_{X_{t-1}} P(X_t|X_{t-1})P(X_{t-1}|Z^{t-1}) \quad (1)$$

The likelihood $P(Z_t|X_t)$ expresses the *measurement model*, the probability that we will observe the measurement Z_t given the state X_t at time t. The *motion model* $P(X_t|X_{t-1})$ predicts the state X_t at time t given the previous state X_{t-1}.

2.2. Particle Filters

In a particle filter [10, 11, 4] we recursively approximate the posterior $P(X_{t-1}|Z^{t-1})$ as a set of N weighted samples $\{X_{t-1}^{(i)}, \pi_{t-1}^{(i)}\}_{i=1}^N$, where $\pi_{t-1}^{(i)}$ is the weight for particle $X_{t-1}^{(i)}$. Given this approximate representation, we obtain a Monte Carlo approximation of the Bayes filtering distribution:

$$P(X_t|Z^t) \approx kP(Z_t|X_t) \sum_i \pi_{t-1}^{(i)} P(X_t|X_{t-1}^{(i)}) \quad (2)$$

One way to view a particle filter is as an importance sampler for this distribution. Specifically, N samples $X_t^{(j)}$ are drawn from the following *proposal distribution* q

$$X_t^{(j)} \sim q(X_t) \triangleq \sum_i \pi_{t-1}^{(i)} P(X_t|X_{t-1}^{(i)})$$

and then weighted by the likelihood, i.e.

$$\pi_t^{(j)} = P(Z_t|X_t^{(j)})$$

This results in a weighted particle approximation $\{X_t^{(j)}, \pi_t^{(j)}\}_{j=1}^N$ for the posterior $P(X_t|Z^t)$ at time t. Note that there are alternative ways to view the particle filter that more easily accommodate other variants (see e.g. [1]), but the mixture proposal view above is well suited for our application.

3. Rao-Blackwellization

When modeling a target's appearance with a subspace model, the state space is augmented with appearance coefficients such that the state $X_t = (l_t, a_t)$ now consists of two parts: the *location* part l_t, which models the position of the target, and the *appearance part* a_t, containing the appearance subspace coefficients. The Bayes filter (1) now becomes:

$$P(l_t, a_t|Z^t) = kP(Z_t|l_t, a_t) \times \int_{l_{t-1}} \int_{a_{t-1}} P(l_t, a_t|l_{t-1}, a_{t-1})P(l_{t-1}, a_{t-1}|Z^{t-1}) \quad (3)$$

In the above we obtain the posterior over the current state by integrating over *both* the location l_{t-1} and appearance coefficients a_{t-1} at the previous time-step $t-1$. If we integrate out the appearance part a_t of the state, we obtain a marginal filter for the location l:

$$P(l_t|Z^t) = k \int_{a_t} P(Z_t|l_t, a_t) \times \int_{l_{t-1}} \int_{a_{t-1}} P(l_t, a_t|l_{t-1}, a_{t-1})P(l_{t-1}, a_{t-1}|Z^{t-1}) \quad (4)$$

In the next section, we will approximate this *exact* marginal filter using a hybrid Monte Carlo method, by sampling over the location l_t but analytically representing the appearance part a_t of the target state.

3.1. The Rao-Blackwellized Particle Filter

In particle filtering, "Rao-Blackwellization" (RB) refers to integrating out part of the state analytically, with the result that the variance of the resulting Rao-Blackwellized particle filter (RBPF) is sharply reduced [16]. For the same level of performance fewer samples will be needed since, intuitively, part of the posterior over the state is calculated exactly or analytically approximated rather than approximated using a more expensive and noisy sample set.

In a Rao-Blackwellized particle filter we approximate the posterior $P(l_{t-1}, a_{t-1}|Z^{t-1})$ over the previous joint state by a set of particles $\{l_{t-1}^{(i)}, w_{t-1}^{(i)}, \alpha_{t-1}^{(i)}(a_{t-1})\}_{i=1}^N$, *each with its own conditional distribution* $\alpha_{t-1}^{(i)}(a_{t-1})$

over the appearance a_{t-1} [16]:

$$P(l_{t-1}, a_{t-1}|Z^{t-1}) = P(l_{t-1}|Z^{t-1})P(a_{t-1}|l_{t-1}, Z^{t-1}) \approx \sum_i w_{t-1}^{(i)} \delta(l_{t-1}^{(i)}) \alpha_{t-1}^{(i)}(a_{t-1}) \tag{5}$$

Formally, $\alpha_{t-1}^{(i)}(a_{t-1})$ is defined as the density on a_{t-1} conditioned on particle i and the measurements Z^{t-1}:

$$\alpha_{t-1}^{(i)}(a_{t-1}) \triangleq P(a_{t-1}|l_{t-1}^{(i)}, Z^{t-1}) \tag{6}$$

Substituting the hybrid approximation (5) into the expression for the marginal filter (4), we obtain after some manipulation the following Monte-Carlo approximation to the exact marginal Bayes filter:

$$P(l_t|Z^t) \approx k \sum_i w_{t-1}^{(i)} \int_{a_t} P(Z_t|l_t, a_t) \times \int_{a_{t-1}} P(a_t|l_t, l_{t-1}^{(i)}, a_{t-1}) P(l_t|l_{t-1}^{(i)}, a_{t-1}) \alpha_{t-1}^{(i)}(a_{t-1}) \tag{7}$$

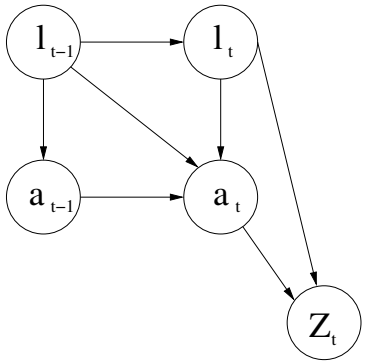

Figure 1: Dynamic Bayes network for the Rao-Blackwellized particle filter. The target state contains both the location l_t and the subspace coefficients a_t

It might seem that we could now implement the same importance sampling scheme as with the particle filter. However, the complicated form of the approximation (7) makes that intractable in general. In theory, it is possible to directly sample from the approximation (and hence do away with the importance weights), but this is both computationally and analytically difficult in all but the simplest cases. Hence, to obtain a practical algorithm we make one additional assumption, namely that the motion model for the location l_t does not depend on the appearance a_{t-1} at time $t-1$:

$$P(l_t|l_{t-1}^{(i)}, a_{t-1}) = P(l_t|l_{t-1}^{(i)}) \tag{8}$$

The dynamic Bayes network for this model is shown in Figure 1. We can now move the motion model out of the integral in (7), obtaining

$$P(l_t|Z^t) \approx k \sum_i w_{t-1}^{(i)} P(l_t|l_{t-1}^{(i)}) \times \int_{a_t} P(Z_t|l_t, a_t) \int_{a_{t-1}} P(a_t|l_t, l_{t-1}^{(i)}, a_{t-1}) \alpha_{t-1}^{(i)}(a_{t-1}) \tag{9}$$

Now we can do importance sampling in the usual way, using the empirical predictive density $\sum_i w_{t-1}^{(i)} P(l_t|l_{t-1}^{(i)})$ as the proposal density. The final algorithm is summarized below.

3.2. RB Particle Filter Summary

For each time step t, starting from the posterior approximation $\{l_{t-1}^{(i)}, w_{t-1}^{(i)}, \alpha_{t-1}^{(i)}(a_{t-1})\}_{i=1}^N$, repeat for $j \in 1..N$:

1. Randomly select a particle $l_{t-1}^{(i)}$ from the previous time step according to the weights $w_{t-1}^{(i)}$.

2. Sample a new particle $\hat{l}_t^{(j)}$ from the chosen model:

$$\hat{l}_t^{(j)} \sim P(l_t|l_{t-1}^{(i)})$$

3. *Extra step*: calculate the posterior density $\alpha_t^{(j)}(a_t)$ on the subspace coefficients a_t:

$$\alpha_t^{(j)}(a_t) = k_t^{(j)} P(Z_t|\hat{l}_t^{(j)}, a_t) \times \int_{a_{t-1}} P(a_t|\hat{l}_t^{(j)}, l_{t-1}^{(i)}, a_{t-1}) \alpha_{t-1}^{(i)}(a_{t-1}) \tag{10}$$

Note that this is where the measurement Z_t is integrated, conditioned on the chosen location $\hat{l}_t^{(j)}$. Also, the integral yields a predictive density on a_t given a move from location $l_{t-1}^{(i)}$ to location $l_t^{(j)}$.

4. Calculate the importance weight $w_t^{(j)}$ as the integral over the *unnormalized* $\alpha_t^{(j)}(a_t)$, i.e.,

$$w_t^{(j)} = 1/k_t^{(j)}$$

where $k_t^{(j)}$ is the normalization constant in (10).

4. RB EigenTracking

4.1. Probabilistic PCA

The integrals in the general Rao-Blackwellized particle filter framework become analytically tractable when we

use the Probabilistic PCA (PPCA) model which is detailed in [20, 17]. In PPCA, a d-dimensional image T is generated according to a factor-analysis model

$$T = \mu + Wa_t + n$$

where μ is the mean of the training image data set, a is a q-dimensional vector of latent variables, i.e., in our case a contains the appearance coefficients, and the noise n is distributed according to an isotropic Gaussian noise model $n \sim \mathcal{N}(0, \sigma^2 I_d)$. The $d \times q$ factor matrix W contains scaled versions of the first q principal components U_q of the data (the eigen-images):

$$W = U_q(\Lambda_q - \sigma^2 I_q)$$

Here Λ_q is a diagonal matrix containing the first q eigenvalues $\lambda_i, 1 \leq i \leq q$, and

$$\sigma^2 = \frac{1}{d-q} \sum_{i=q+1}^{d} \lambda_i$$

is the average of the "lost variance" in the remaining $d-q$ principal directions. We define dependence on the sampled location $\hat{l}_t^{(j)}$ through an inverse warp R^{-1} on the image measurement Z_t. The warp obtains an image vector $T^{(j)} = R^{-1}(Z_t, \hat{l}_t^{(j)})$ from a rectangular region with an offset center of rotation with position and orientation $\hat{l}_t^{(j)}$. Taking all this together, we obtain the following Gaussian likelihood model:

$$P(Z_t|\hat{l}_t^{(j)}, a_t) = \frac{1}{\sqrt{|2\pi I_d \sigma^2|}} \times \exp -\frac{1}{2}\|T^{(j)} - \mu - Wa_t\|^2_{I_d\sigma^2} \quad (11)$$

where we write the squared Mahalanobis distance as

$$\|t-\mu\|^2_\Sigma \triangleq (t-\mu)^T\Sigma^{-1}(t-\mu)$$

4.2. "Extra Step": Analytical Update

The posterior density $\alpha_t(a_t)$ on the current appearance coefficients in (10) can be shown to be a normal density $\mathcal{N}(a_t; \hat{a}_t, P_t)$ by assuming that the coefficients a_t change smoothly over time according to a Gaussian process:

$$a_t|\hat{l}_t^{(j)}, l_{t-1}^{(i)}, a_{t-1} \sim \mathcal{N}(a_{t-1}, Q_{ij})$$

The diagonal variance Q_{ij} can be adapted to the distance between the locations $\hat{l}_t^{(j)}$ and $l_{t-1}^{(i)}$, useful in many domains where the target appearance tends to change more when the target moves quickly. Using the PPCA image model (11), and under the inductive assumption that the density $\alpha_{t-1}(a_{t-1})$ on the previous appearance coefficients a_{t-1} is a normal density $\mathcal{N}(a_{t-1}; \hat{a}_{t-1}, P_{t-1})$, we obtain the following Kalman filter update equations for the subspace coefficients:

$$\begin{aligned} P_t^{(j)} &= \left(\sigma^{-2}W^TW + (Q_{ij} + P_{t-1})^{-1}\right)^{-1} \quad (12) \\ \hat{a}_t^{(j)} &= P_t \times (\sigma^{-2}W^Ty + (Q_{ij} + P_{t-1})^{-1}\hat{a}_{t-1}) \end{aligned}$$

where $y \triangleq T^{(j)} - \mu$.

4.3. Calculating the Importance Weights

The importance weight can be derived by leveraging the Gaussian assumptions made in the measurement model and analytical update. The result of the integral over the previous appearance coefficients a_{t-1} in (10) is a Gaussian $\mathcal{N}(a_{t-1}; \hat{a}_{t-1}, Q_{ij}+P_{t-1})$. Since the measurement model in (11) is also Gaussian, we compute the importance weight according to

$$\begin{aligned} w_t^{(j)} &= k \exp -\frac{1}{2}\Big[\|T^{(j)} - \mu - W\hat{a}_t\|^2_{I_d\sigma^2} + \\ &\quad \|\hat{a}_t - \hat{a}_{t-1}\|^2_{Q_{ij}+P_{t-1}}\Big] \quad (13) \end{aligned}$$

and

$$k = \frac{\sqrt{|2\pi P_t|}}{\sqrt{|2\pi I_d\sigma^2|}\sqrt{|2\pi(Q_{ij}+P_{t-1})|}}$$

Intuitively, a small importance weight is given to particles with high reconstruction error using the maximum a posteriori (MAP) estimates of the current appearance coefficients as well as particles where the difference between the MAP estimates of the coefficients differ greatly between time steps.

4.4. EigenTracking Algorithm Summary

For convenience, we include a complete summary of the tracking algorithm. We track target position and orientation such at $l_t = [x, y, \theta]^T$. Starting from the posterior approximation at time t, a set of weighted hybrid particles $\{l_{t-1}^{(i)}, w_{t-1}^{(i)}, \hat{a}_{t-1}^{(i)}, P_{t-1}^{(i)}\}_{i=1}^N$, repeat for $j \in 1..N$:

1. Randomly select a particle $l_{t-1}^{(i)}$ from the previous time step according to the weights $w_{t-1}^{(i)}$.
2. Sample from the motion model $P(l_t|l_{t-1}^{(i)})$ for the chosen particle to obtain a predicted location $\hat{l}_t^{(j)}$.
3. Update $\hat{a}_{t-1}^{(i)}$ and $P_{t-1}^{(i)}$ to $\hat{a}_t^{(j)}$ and $P_t^{(j)}$ according to the filtering equations in (12).
4. Calculate the weight $w_t^{(j)}$ according to (13).

Repeating this N times yields the weighted hybrid particle set $\{\hat{l}_t^{(j)}, w_t^{(j)}, \hat{a}_t^{(j)}, P_t^{(j)}\}_{j=1}^N$ at the current time step t, approximating the posterior $P(l_t, a_t|Z^t)$.

5. Experimental Results

We used the Rao-Blackwellized particle filter to track an unmarked honey bee in an observation hive. The automated recording of trajectory data has significant applications in the study of honey bee behavior and physiology [19]. This application presents substantial challenges due to temporary occlusions that occur in the close proximity of other bees, complex variations in the appearance of the tracked bee, and the unpredictability of the tracked bee's movements.

(a)

(b)

Figure 2: (a) Tracking honey bees in a hive presents substantial challenges for a tracker due to temporary occlusions that occur in the close proximity of other bees and complex variations in the appearance of the tracked bee. (b) Mean target image with the first 5 principal components, or eigen-bees, estimated from a training set

The test sequence was recorded at 15fps at 720x480 pixels and scaled to 360x240 pixels (see Figure 2a). We used 40 iterations of EM (see [20]) to learn the PPCA image model from a training image data set consisted of 172 color bee images, each measuring 14 by 31 pixels. The mean image and the first five principal components (eigen-bees) are shown in Figure 2b. The center of rotation of the rectangle was offset 15 pixels from the top of the image. Because the motion of the bee is difficult to predict, we used a Gaussian motion model where the target state is updated according to $l|l_{t-1} = A(\theta_{t-1})\Delta l + l_{t-1}$ by sampling from a zero mean Gaussian $\Delta l \sim N(0, \text{diag}(\sigma^2_{\Delta x'}\sigma^2_{\Delta y'}\sigma^2_{\Delta\theta'}))$ with preset variances, $\sigma_{\Delta x} = 2$, $\sigma_{\Delta y} = 3$, and $\sigma_{\Delta\theta} = 0.3$ for position and orientation, where A is matrix that rotates a location l according to the angle θ_{t-1}. The variance Q_{ij} in the coefficients was set to and $7I_q$ and $10I_q$ when the position of the target changed by more than 2 pixels.

Components	Failures	MSE
0	11	14.71 ± 34.73
3	8	13.47 ± 32.61
6	4	7.85 ± 18.80
9	1	6.43 ± 14.89
12	0	4.73 ± 10.47

Table 1: When using the RBPF with 500 hybrid particles, increasing the complexity of the PPCA model by adding principal components both improved the quality of the trajectory and decreased the number of failures.

Particles	Failures	MSE
500	69	52.76 ± 67.68
1500	44	42.37 ± 63.19
3000	33	31.71 ± 51.10
6000	20	21.11 ± 45.03

Table 2: In a "plain" particle filter where the state vector contains 12 appearance coefficients which are sampled, instead of computed analytically, the number of particles required was increased substantially.

Using increasingly rich subspace models by adding additional PPCA components both decreased the number of failures and improved track quality. In Table 1 we show the results of using the RBPF with 500 hybrid particles. A failure was recorded when the position reported tracker deviated more than half the width of the target, 7 pixels in the image, from the ground truth trajectory. When a failure occurred, the tracker was reinitialized at the ground truth position and allowed to resume tracking. We measured the quality of the trajectory by computing the mean squared error (MSE), or squared distance, of the position reported by tracker from the ground truth. The zero principal component model is equivalent to using a simple Gaussian image model with mean μ and variance $\sigma^2 I_d$, which results in poor tracking. With q=12 components in the PPCA model we report no failures and track quality markedly improved.

The same results are qualitatively illustrated in Figure 3, where we have shown the corresponding 2D trajectories for the $q = 0$ and $q = 12$ component models. Through the correlations between pixel color captured by the principal components the tracker is able to more

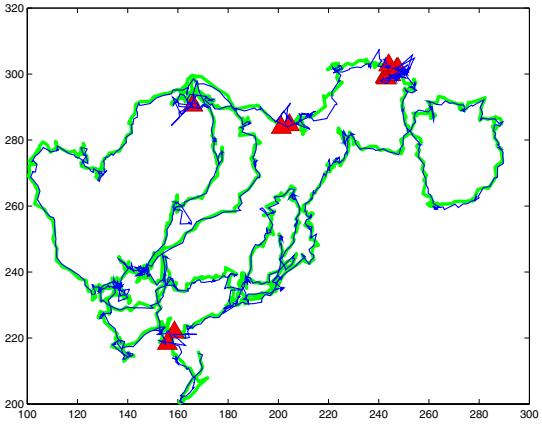

(a) $q = 0$, using 500 hybrid particles

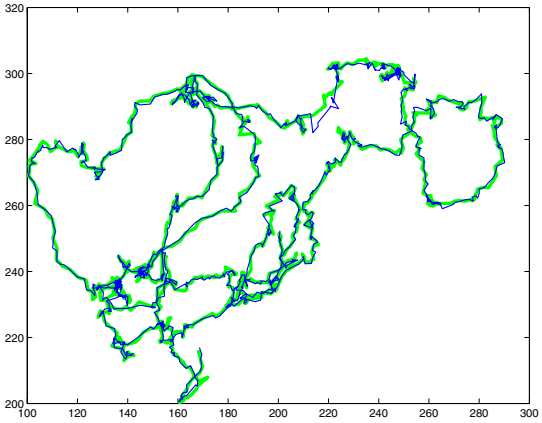

(b) $q = 12$, using 500 hybrid particles

Figure 3: As the complexity of the PPCA image model is increased the tracking quality improves. The blue recorded trajectory is shown superimposed on the green ground truth trajectory. Failures are designated by red triangles. In (a) and (b) we show trajectories obtained by models with zero and 12 components respectively.

accurately the model the appearance of the target.

In contrast to the RBPF results, the number of particles required in a standard particle filter to track at a worse level of performance grows intractably large. Table 2 shows the results we obtained using q=12, for increasing numbers of particles. In Figure 4, we show the trajectory for standard particle filtering, using 500 particles to enable comparison with Figure 3. The analytical update of appearance coefficients in the RB particle filter substantially decreases the number of samples required to track a target. For a standard particle filter with 500 particles and 12 principal components 69 failures were recorded. Whereas, zero failures were recorded for the RB filter with 500 hybrid particles. Adding particles to the standard particle filter does improve performance,

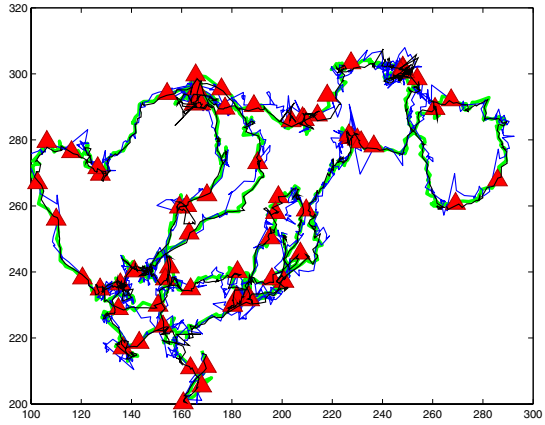

$q = 12$, using 500 particles

Figure 4: In a standard particle filter, the number of particles required increases exponentially as the state vector contains 15 dimensions: the location, orientation, and 12 appearance coefficients. For comparison with Figure 3 in which we use an RB filter with 500 hybrid particles, we plot the trajectory obtained by a standard particle filter with 500 particles. The blue recorded trajectory is shown superimposed on the green ground truth trajectory. Failures are designated by red triangles.

but not nearly enough.

This result is not surprising, considering the fact that the state space sampled over is 15-dimensional: 3 degrees of freedom for the the location and orientation, and 12 additional subspace coefficients to model appearance. It is well known that importance sampling, which lies at the core of the particle filtering approach (see Section 2.2), breaks down in high-dimensional spaces.

6. Conclusions

In this paper, we introduced an efficient method for using subspace representations as the measurement model in a particle filter. We applied Rao-Blackwellization to integrate out the subspace coefficients of the state vector, and have shown how this can be done in a general Rao-Blackwellized particle filter framework. For an application that is of interest to researchers who study animal behavior, we specialized the RBPF framework to visual tracking in a way that closely mimics Jepson and Black's influential work on the EigenTracking algorithm. One possible downside of using a Gaussian latent variable model is that outliers in the image noise, such as specular reflections, are not well modeled using a Gaussian noise model. Hence, one avenue for future work is examining the use of robust, non-Gaussian subspace models in the context of the RBPF.

We tested the algorithm in a challenging tracking application in which we tracked a target in clutter. The conclusions we attach to our experiments are twofold:

- Using a subspace appearance model in conjunction with a particle filter can substantially increase tracking performance and decrease the number of tracking failures, especially when tracking targets that exhibit complex appearance changes over time.
- The analytical update of the subspace coefficients in the RBPF sharply reduces the number of particles needed and hence the computational requirements to accurately track a target in clutter.

Based on these promising results, we expect that hybrid sampling approaches will play an increasingly important role in computer vision tracking applications.

Acknowledgments

This work was funded under NSF Award IIS-0219850.

References

[1] S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for on-line non-linear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50(2):174–188, February 2002.

[2] P. N Belhumeur, J P Hespanha, and David J Kriegman. Eigenfaces vs. fisherfaces: Recognition using class specific linear projection. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1996.

[3] M.J. Black and A.D. Jepson. Eigentracking: Robust matching and tracking of articulated objects using a view-based representation. In *Eur. Conf. on Computer Vision (ECCV)*, 1996.

[4] J. Carpenter, P. Clifford, and P. Fernhead. An improved particle filter for non-linear problems. Technical report, Department of Statistics, University of Oxford, 1997.

[5] G. Casella and C.P. Robert. Rao-Blackwellisation of sampling schemes. *Biometrika*, 83(1):81–94, March 1996.

[6] T.F. Cootes, Edwards G, J, and C.J. Taylor. Active appearance models. In *Eur. Conf. on Computer Vision (ECCV)*, 1998.

[7] T.F. Cootes, A. Hill, C.J. Taylor, and J. Haslam. The use of active shape models for locating structures in medical images. *Image and Vision Computing*, 12(6):355–366, July 1994.

[8] G.J. Edwards, T.F. Cootes, and C.J. Taylor. Face recognition using active appearance models. In *Eur. Conf. on Computer Vision (ECCV)*, 1998.

[9] N. Freitas. Rao-Blackwellised particle filtering for fault diagnosis. *IEEE Aerospace*, 2002.

[10] N.J. Gordon, D.J. Salmond, and A.F.M. Smith. Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *IEE Procedings F*, 140(2):107–113, 1993.

[11] M. Isard and A. Blake. Contour tracking by stochastic propagation of conditional density. In *Eur. Conf. on Computer Vision (ECCV)*, pages 343–356, 1996.

[12] M. Isard and A. Blake. Condensation – conditional density propagation for visual tracking. *Intl. J. of Computer Vision*, 29(1):5–28, 1998.

[13] Z. Khan, T. Balch, and F. Dellaert. Efficient particle filter-based tracking of multiple interacting targets using an MRF-based motion model. In *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Las Vegas, 2003.

[14] Yang M., Kriegman D., and Ahuja N. Detecting faces in images: A survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(1):34–58, January 2002.

[15] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. FastSLAM: A factored solution to the simultaneous localization and mapping problem. In *AAAI Nat. Conf. on Artificial Intelligence*, 2002.

[16] K. Murphy and S. Russell. Rao-Blackwellised particle filtering for dynamic bayesian networks. In A. Doucet, N. de Freitas, and N. Gordon, editors, *Sequential Monte Carlo Methods in Practice*. Springer-Verlag, New York, January 2001.

[17] S. Roweis. EM algorithms for PCA and SPCA. In Michael I. Jordan, Michael J. Kearns, and Sara A. Solla, editors, *Advances in Neural Information Processing Systems*, volume 10. The MIT Press, 1998.

[18] D. Schulz, D. Fox, and J. Hightower. People tracking with anonymous and id-sensors using rao-blackwellised particle filters. In *Intl. Joint Conf. on Artificial Intelligence (IJCAI)*, 2003.

[19] T.D. Seeley. *The Wisdom of the Hive*. Harvard University Press, 1995.

[20] M.E. Tipping and C.M. Bishop. Probabilistic principal component analysis. Technical Report NCRG/97/010, Neural Computing Research Group, Aston University, September, 1997.

[21] M. Turk and A.P. Pentland. Face recognition using eigenfaces. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 586–591, 1991.

[22] T. Vetter and V. Blanz. Coloured 3d face models from single images: An example based approach. In *Eur. Conf. on Computer Vision (ECCV)*, pages 499–513, 1998.