

OPTIMAL POSTERIOR DENSITY APPROXIMATION FOR BEARINGS-ONLY TRACKING

Mark R. Morelande¹, Lennart Svensson², Jonas Hagmar³, Mats Jirstrand³

¹Dept. of Electrical and Electronic Engineering, The University of Melbourne

²Dept. of Signals and Systems, Chalmers University of Technology

³Dept. of System Biology and Bioimaging, Fraunhofer-Chalmers Centre

ABSTRACT

The optimal homotopy filter is a nonlinear filtering approximation which seeks an optimal parameterisation for the posterior. The search for an optimal parameterisation is performed by constructing a homotopy between the prior and posterior and solving the resulting ordinary differential equation. Here the optimal homotopy filter is applied to the problem of bearings-only tracking. A simulation analysis shows that the performance of the optimal homotopy filter compares favourably to established algorithms.

Index Terms— Nonlinear filtering; Bayesian estimation; bearings-only tracking;

1. INTRODUCTION

The key quantity in nonlinear filtering is the posterior density from which posterior statistics of interest can be computed. Of particular interest in point estimation is the posterior mean since this is the minimum mean square error estimator. It is only rarely possible to compute the posterior density exactly. A well-known example where exact computation is possible is when both the dynamic and measurement equations are linear and Gaussian. In this case the posterior is Gaussian and can be computed using the Kalman filter (KF). More generally, for non-Gaussian/nonlinear dynamic and/or measurement equations the posterior is usually not available in closed-form and must be approximated.

A common method is to adopt a Gaussian approximation to the posterior with the mean and covariance matrix computed by an approximation of the KF recursion. There are many variants of this approach including the extended KF (EKF) [1] and the unscented KF (UKF) [2]. These algorithms have the desirable property of computational simplicity but are not suitable for all problems. Sequential Monte Carlo methods, or particle filters (PFs), are of more general applicability. A PF approximates the posterior by a random sample, usually drawn from a suitable importance density [3]. In principle, PFs can be used to provide accurate approximation of the posterior in a large class of problems. However, in practice sufficiently accurate results often require a prohibitively large sample size, even with sophisticated sampling strategies.

The approach to nonlinear filtering taken here is to seek an optimal approximation, in some sense, to the posterior from within a certain class of distributions. The posterior is represented by a finite number of parameters which are updated as measurements are acquired by defining a homotopy between the prior and the posterior, i.e., a continuous transformation from the prior to the posterior as a scalar parameter is varied between 0 and 1. The parameters of the posterior approximation are then found by solving an ordinary differential equation (ODE). This idea, which we refer to as optimal homotopy, has previously been proposed in [4, 5]. Closely related

to optimal homotopy is the idea of particle flow proposed in [6]. As in a PF, particle flow filters represent the posterior by a collection of points. However, instead of random sampling, particle flow methods construct an ODE to adjust the location of each point. The particle flow methods are therefore similar in spirit to optimal homotopy although the principles used to construct the ODEs are quite different for the two methods.

While promising, the performance benefits of optimal homotopy have not been established as it has not yet been implemented as a filtering recursion. The main contribution of this paper is the application of optimal homotopy to the problem of tracking an object moving in two dimensions using bearings measurements. An interesting property of this nonlinear filtering problem is that full observability of the state is subject to the observer motion. For example, full observability of the state of an object moving with constant velocity requires the observer to perform a manoeuvre. The main impediments to the application of optimal homotopy to bearings-only tracking (BOT) are the solution of the ODE describing the evolution of the parameters of the posterior density approximation and the need to evaluate several intractable integrals. Several approaches to these problems are developed and their performances are analysed and compared with existing algorithms.

The paper is organised as follows. The optimal homotopy filter is described in Section 2 and applied to BOT in Section 3. A simulation analysis, including comparisons with existing algorithms, is given in Section 4.

2. THE OPTIMAL HOMOTOPY FILTER

In this section the general idea of optimal homotopy is reviewed. We consider a stochastic dynamic system with state vector $\mathbf{x}_k \in \mathbb{R}^c$ which evolves as

$$\mathbf{x}_k | \mathbf{x}_{k-1} \sim v(\cdot | \mathbf{x}_{k-1}) \quad (1)$$

where $v(\cdot | \cdot)$ is the transition density. The state is observed through measurements $\mathbf{y}_k \in \mathbb{R}^d$ generated according to

$$\mathbf{y}_k | \mathbf{x}_k \sim g(\cdot | \mathbf{x}_k) \quad (2)$$

Let $\pi_k(\cdot)$ denote the posterior density of the state at time k , i.e., the density of \mathbf{x}_k conditional on the measurements $\mathbf{y}_1, \dots, \mathbf{y}_k$. The posterior can be computed recursively using Bayes' rule:

$$\pi_k(\mathbf{x}_k) \propto g(\mathbf{y}_k | \mathbf{x}_k) \varpi_k(\mathbf{x}_k) \quad (3)$$

where $\varpi_k(\cdot)$ is the prior density,

$$\varpi_k(\mathbf{x}_k) = \int v(\mathbf{x}_k | \mathbf{x}_{k-1}) \pi_{k-1}(\mathbf{x}_{k-1}) d\mathbf{x}_{k-1} \quad (4)$$

Eq. (3) is commonly referred to as the correction, or update, step while (4) is usually called the prediction step. The object of nonlinear filtering is to accurately approximate the recursion consisting of both the prediction and correction steps.

Let $f(\cdot; \theta)$ denote a family of densities parameterised by the r -dimensional vector θ . It is desired to select θ so that $f(\cdot; \theta)$ approximates the posterior as closely as possible. The optimal homotopy filter (OHF) focuses on the correction step (3), i.e., assuming a parameterisation for the prior is available, the OHF seeks an optimal parameterisation for the posterior. The motivation for the focus on the correction step is that this part of the recursion is potentially a much greater source of approximation error than the prediction step [6, 7].

Let $f(\cdot; \theta_{k|k-1})$ denote the OHF approximation to the prior. Replacing the true prior in (3) with the OHF approximation gives the unnormalised posterior density approximation

$$\tilde{\pi}_k(\mathbf{x}_k) = g(\mathbf{y}_k | \mathbf{x}_k) f(\mathbf{x}_k; \theta_{k|k-1}) \quad (5)$$

Note that $\tilde{\pi}_k(\cdot)$ will not, in general, belong to the same class of densities as the prior approximation $f(\cdot; \theta_{k|k-1})$, i.e., there is no θ such that $\tilde{\pi}_k(\mathbf{x}) \propto f(\mathbf{x}; \theta)$. Instead we select a value of θ for which $f(\cdot; \theta)$ matches $\tilde{\pi}_k(\cdot)$ as closely as possible. Let $J(\cdot, \cdot)$ denote a functional, such as the Kullback-Leibler divergence (KLD) or the integrated squared difference, which measures the difference between two densities. Then, the OHF approximates the posterior density at time k by $f(\cdot; \theta_{k|k})$ where

$$\theta_{k|k} = \underset{\theta}{\operatorname{argmin}} J(\tilde{\pi}_k(\cdot), f(\cdot; \theta)) \quad (6)$$

The optimisation (6) is difficult to perform in general. The OHF solves this difficulty by constructing a homotopy between the prior approximation $f(\cdot; \theta_{k|k-1})$ and the unnormalised posterior approximation $\tilde{\pi}_k(\cdot)$. The particular homotopy used here is as follows, for $0 \leq \lambda \leq 1$,

$$\tilde{\pi}_k(\mathbf{x}_k; \lambda) = g(\mathbf{y}_k | \mathbf{x}_k)^\lambda f(\mathbf{x}_k; \theta_{k|k-1}) \quad (7)$$

Let

$$\theta_k(\lambda) = \underset{\theta}{\operatorname{argmin}} J(\tilde{\pi}_k(\cdot; \lambda), f(\cdot; \theta)) \quad (8)$$

Note that $\theta_k(0) = \theta_{k|k-1}$ while $\theta_k(1) = \theta_{k|k}$ is the desired solution of (6). It will be convenient to denote the criterion for the difference between $\tilde{\pi}_k(\cdot; \lambda)$ and $f(\cdot; \theta)$ as $J(\lambda, \theta)$. We require that

$$\nabla_{\theta} J(\lambda, \theta)|_{\theta=\theta_k(\lambda)} = \mathbf{0} \quad (9)$$

where $\nabla_{\theta} = [\partial/\partial\theta_1, \dots, \partial/\partial\theta_r]'$. It follows from (9) that $\nabla_{\theta} J(\lambda, \theta)|_{\theta=\theta_k(\lambda)}$ remains constant as λ varies between 0 and 1. Therefore, differentiating the LHS of (9) with respect to λ gives

$$\mathbf{H}(\lambda, \theta_k) \frac{d\theta_k(\lambda)}{d\lambda} + d(\lambda, \theta_k(\lambda)) = \mathbf{0} \quad (10)$$

where $\mathbf{H}(\lambda, \theta_k(\lambda)) = \nabla_{\theta} \nabla'_{\theta} J(\lambda, \theta)|_{\theta=\theta_k(\lambda)}$ and

$$d(\lambda, \theta_k(\lambda)) = \nabla_{\lambda} \nabla_{\theta} J(\lambda, \theta)|_{\theta=\theta_k(\lambda)} \quad (11)$$

This leads to the following ODE for the parameter vector:

$$\frac{d\theta_k(\lambda)}{d\lambda} = -\mathbf{H}(\lambda, \theta_k(\lambda))^{-1} d(\lambda, \theta_k(\lambda)) \quad (12)$$

with initial condition $\theta_k(0) = \theta_{k|k-1}$. Solution of the first-order ODE (12) provides the optimal parameterisation of the posterior at time k . We consider two approximate ways of solving (12).

The first method of approximating (12) is a simple Euler scheme. Let $0 = \lambda_0 < \lambda_1 < \dots < \lambda_m = 1$. Then, for $i = 0, \dots, m-1$, we obtain the homotopy difference equation, referred to by the acronym H Δ E in [5], as

$$\theta_k(\lambda_{i+1}) = \theta_k(\lambda_i) - s_{i+1} \mathbf{H}(\lambda_i, \theta_k(\lambda_i))^{-1} d(\lambda_i, \theta_k(\lambda_i)) \quad (13)$$

where $s_i = \lambda_i - \lambda_{i-1}$. The solution (13) becomes exact as the step sizes $s_i \rightarrow 0$. However, numerical implementation of (13) can be quite difficult, as will be seen in Section 3 where we consider the application of the OHF to BOT. The second, simpler way of approximately solving (12) is to replace $\tilde{\pi}_k(\cdot; \lambda)$ with, for $\lambda_i \leq \lambda \leq \lambda_{i+1}$,

$$\tilde{\pi}_{k,i}(\mathbf{x}_k; \lambda) = g(\mathbf{y}_k | \mathbf{x}_k)^{\lambda - \lambda_i} f(\mathbf{x}_k; \theta_k(\lambda_i)) \quad (14)$$

Let $\tilde{J}(\lambda, \theta) = J(\tilde{\pi}_{k,i}(\cdot; \lambda), f(\cdot; \theta))$ and

$$\tilde{\mathbf{H}}(\lambda, \theta_k(\lambda)) = \nabla_{\theta} \nabla'_{\theta} \tilde{J}(\lambda, \theta)|_{\theta=\theta_k(\lambda)}, \quad (15)$$

$$\tilde{d}(\lambda, \theta_k(\lambda)) = \nabla_{\lambda} \nabla_{\theta} \tilde{J}(\lambda, \theta)|_{\theta=\theta_k(\lambda)} \quad (16)$$

Then the approximate H Δ E (AH Δ E) is

$$\theta_k(\lambda_{i+1}) = \theta_k(\lambda_i) - s_{i+1} \tilde{\mathbf{H}}(\lambda_i, \theta_k(\lambda_i))^{-1} \tilde{d}(\lambda_i, \theta_k(\lambda_i)) \quad (17)$$

Unlike the H Δ E of (13), the AH Δ E (17) does not become an exact solution of (12) as the step sizes $s_i \rightarrow 0$ because the densities (5) and (14) are generally not equal. However, the AH Δ E is often much simpler to implement than the H Δ E. This is the case for the application of the OHF to BOT presented in the following section.

3. APPLICATION OF OPTIMAL HOMOTOPY TO BEARINGS-ONLY TRACKING

In BOT the state is given by $\mathbf{x}_k = [x_k, y_k, \dot{x}_k, \dot{y}_k]'$ where (x_k, y_k) is the object position in Cartesian coordinates and the dot notation indicates differentiation with respect to time. The motion model for the object includes random perturbations in velocity to model small object manoeuvres. This results in the dynamic model

$$v(\mathbf{x}_k | \mathbf{x}_{k-1}) = N(\mathbf{x}_k; \mathbf{F} \mathbf{x}_{k-1}, \mathbf{Q}) \quad (18)$$

where $N(\cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ is the Gaussian density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ and

$$\mathbf{F} = \begin{bmatrix} 1 & T \\ 0 & 1 \end{bmatrix} \otimes \mathbf{I}_2 \quad (19)$$

$$\mathbf{Q} = q \begin{bmatrix} T^3/3 & T^2/2 \\ T^2/2 & T \end{bmatrix} \otimes \mathbf{I}_2 \quad (20)$$

In (19) and (20), $q > 0$ is the process noise intensity, T is the measurement sampling period, \otimes is the Kronecker product and \mathbf{I}_m is the $m \times m$ identity matrix.

At the k th sampling instant, a measurement y_k of the object bearing is made from an observer located at $\boldsymbol{\rho}_k = (\xi_k, \zeta_k)$ in Cartesian coordinates. The measurement satisfies

$$g(y_k | \mathbf{x}_k) = N(y_k; h(\mathbf{x}_k, \boldsymbol{\rho}_k), \sigma^2) \quad (21)$$

where

$$h(\mathbf{x}_k, \boldsymbol{\rho}_k) = \arctan \left(\frac{y_k - \zeta_k}{x_k - \xi_k} \right) \quad (22)$$

The first step in the development of an OHF for BOT is to choose a parameterisation. We use a Gaussian distribution due to its simplicity and because it is capable of providing a reasonably accurate

approximation to the true posterior. Another benefit of using a Gaussian approximation to the posterior is that the prediction step (4) can be performed in closed-form for the transition density (18).

Since the measurement depends only on the object position, the homotopy need only be performed over the position elements of the object state. Thus given a Gaussian approximation to the prior, we solve the homotopy differential equation (12), or rather one of the approximate difference equations (13) or (17), to find an optimal parameterisation for the posterior of the position elements. Some simple manipulations can then be used to find the joint posterior for the entire object state. The homotopy parameter vector θ , which represents the Gaussian approximation to the posterior of the object position, is then composed of $r = 5$ elements: the two position means and variances and the covariance between the positions. The position elements of the state vector are collected into \mathbf{p}_k .

The second step is to choose a criterion for measuring the difference between the posterior and the parameterised approximation. We use the KLD which, for densities $f_1(\cdot)$ and $f_2(\cdot)$, is defined as

$$J(f_1(\cdot), f_2(\cdot)) = \int \log \left(\frac{f_1(\mathbf{p})}{f_2(\mathbf{p})} \right) f_1(\mathbf{p}) d\mathbf{p} \quad (23)$$

Although we would like to measure the difference between the posterior and the parameterised approximation only the unnormalised posterior (7) is available. An important property of the KLD is that the use of the unnormalised posterior does not affect the selection of an optimal parameterisation. To see this, let $f_1(\mathbf{p}) = \pi_k(\mathbf{p}) = c\tilde{\pi}_k(\mathbf{p})$ where c is a normalisation factor and $f_2(\mathbf{p}) = f(\mathbf{p}; \theta)$. Substituting into (23) gives

$$J(\pi_k(\cdot), f(\cdot; \theta)) = \log(c) + cJ(\tilde{\pi}_k(\cdot), f(\cdot; \theta)) \quad (24)$$

It can be seen from (24) that the same value of θ which minimises $J(\tilde{\pi}_k(\cdot), f(\cdot; \theta))$ also minimises $J(\pi_k(\cdot), f(\cdot; \theta))$. The need for a criterion which can properly handle comparison with an unnormalised posterior has not been recognised previously.

The key difficulty now is evaluating the gradient matrices and vectors which appear in the difference equations (13) and (17). We consider the quantities required for the H Δ E (13) first. The Hessian matrix is

$$\mathbf{H}(\lambda, \theta_k(\lambda)) = \int \left[\nabla_{\theta} \log(f(\mathbf{p}; \theta)) \nabla_{\theta'} \log(f(\mathbf{p}; \theta)) - \frac{\nabla_{\theta} \nabla_{\theta'} f(\mathbf{p}; \theta)}{f(\mathbf{p}; \theta)} \right] \bigg|_{\theta=\theta_k(\lambda)} \tilde{\pi}_k(\mathbf{p}; \lambda) d\mathbf{p} \quad (25)$$

$$\mathbf{d}(\lambda, \theta_k(\lambda)) = \int \nabla_{\theta} \log(f(\mathbf{p}; \theta)) \bigg|_{\theta=\theta_k(\lambda)} \times \log(g(y_k|\mathbf{p})) \tilde{\pi}_k(\mathbf{p}; \lambda) d\mathbf{p} \quad (26)$$

The exact form of the derivatives has been omitted for the sake of brevity. Neither of the integrals (25), (26) can be evaluated in closed form so approximations are required. After investigation of numerous methods, including Monte Carlo approximations and Taylor series expansions, we found that a change of variables to polar coordinates followed by a grid approximation produces reasonably accurate approximation with moderate computational expense.

For the AH Δ E (17), the required quantities are, for $\lambda_i \leq \lambda \leq$

λ_{i+1} ,

$$\tilde{\mathbf{H}}(\lambda, \theta_k(\lambda)) = \int \left[\nabla_{\theta} \log(f(\mathbf{p}; \theta)) \nabla_{\theta'} \log(f(\mathbf{p}; \theta)) - \frac{\nabla_{\theta} \nabla_{\theta'} f(\mathbf{p}; \theta)}{f(\mathbf{p}; \theta)} \right] \bigg|_{\theta=\theta_k(\lambda)} f(\mathbf{p}; \theta_k(\lambda_i)) d\mathbf{p} \quad (27)$$

$$\tilde{\mathbf{d}}(\lambda, \theta_k(\lambda)) = \int \nabla_{\theta} \log(f(\mathbf{p}; \theta)) \bigg|_{\theta=\theta_k(\lambda)} \times \log(g(y_k|\mathbf{p})) f(\mathbf{p}; \theta_k(\lambda_i)) d\mathbf{p} \quad (28)$$

Eq. (27) can be evaluated in closed-form. This is an important benefit of the AH Δ E compared to the H Δ E. Also, although (28) is intractable, it has a simpler form than the corresponding equation for the H Δ E. The same approximation technique used for (25) and (26) is used for (28).

4. SIMULATION RESULTS

The simulation scenario used to assess the performance of the OHF is shown in Fig. 1. In this scenario the object begins at a range of 5000 m and bearing 80°. The object initially moves with speed 123.5 m/s and heading -130°. After that the object state evolves as in (18) with process noise intensity $q = 0.01$. The observer starts at the origin of the coordinate frame and moves with a constant speed of 154 m/s. The initial heading is -45°. Between the 13th and 17th sampling instants the observer executes a coordinated turn with turn rate 30°/s. It is only after this manoeuvre that the object state becomes fully observable. The sampling period is $T = 1$ s.




Fig. 1. Simulation scenario: the solid line is the observer trajectory and the dashed line is the nominal object trajectory. Crosses indicate the starting positions and solid circles indicate sampling instants.

The OHF is implemented with both the H Δ E and AH Δ E. The former is referred to as simply the OHF while the latter is called the approximate OHF (AOHF). Comparisons are performed with two benchmark existing algorithms. The first is the UKF of [2] implemented in log polar coordinates (LPC) [8]. The UKF-LPC is regarded as one of the most accurate Gaussian approximations. The second algorithm is a PF with the importance density a linearised approximation to the optimal importance density [3]. This is a reasonably sophisticated PF which attempts to use measurement information to direct particles efficiently. The PF is implemented with a sample size of 100 000.

The algorithms are initialised using the first measurement and prior information about the range, speed and heading of the object. Specifically, the prior range distribution is $N(8000, 3500^2)$,

the prior speed distribution is $N(150, 50^2)$ and the prior heading distribution is, in radians, $N(-5\pi/9, \pi^2/12)$. Note that the prior distributions are not particularly informative.

4.1. Demonstration of the OHF

Before comparing the algorithms' performances we demonstrate the operation of the OHF. We consider the 17th sampling instant because the measurement received at the end of the observer manoeuvre results in a dramatic increase in information. This greatly increases tracking accuracy but is challenging for filtering approximations. Fig. 2 shows the error ellipses of the optimal parameterisation found by the OHF for $\lambda = i/4$, $i = 0, \dots, 4$ as well as the object position. The OHF seems to perform well in this instance with the error ellipses gradually concentrating around the true object position as λ increases.




Fig. 2. Demonstration of the OHF: error ellipses for the optimal parameterisation at the 17th sampling instant for $\lambda = 0$ (magenta), $1/4$ (blue), $1/2$ (red), $3/4$ (green) and 1 (black). The true object position is indicated by a cross.

4.2. Performance analysis

Performance statistics for the various algorithms are found by averaging over 500 realisations of the simulation scenario. The RMS position errors are plotted in Fig. 3. The best performance by a significant margin, in terms of final tracking error, is achieved by the OHF. Both the AOHF and OHF are noticeably superior to the UKF-LPC and the PF in this example. It may be expected that the OHF would outperform the AOHF since it attempts an exact solution to the homotopy differential equation (12). The situation is complicated somewhat by the fact that approximation of the Hessian matrix (25) and gradient vector (26) is more difficult in the OHF compared to the AOHF. Indeed the Hessian is available in closed-form for the AOHF. Thus, it is possible that larger errors introduced in the approximation of (25) and (26) contribute to the slightly slower reduction in error of the OHF at the onset of the observer manoeuvre. The AOHF has about a quarter of the computational expense of the OHF. Both algorithms are much more expensive than the UKF-LPC but considerably less expensive than the PF for the sample size of 100 000 used here.

5. CONCLUSIONS

The optimal homotopy filter was applied to the challenging problem of bearings-only tracking. Two variants of the optimal homotopy




Fig. 3. RMS position errors for the UKF-LPC (black), PF (green), OHF (blue) and AOHF (red) plotted against time.

filter were implemented and shown to produce much better performance in a simulation example than existing algorithms. Importantly, the proposed filters achieve this high level of performance with moderate computational expense. Although the strong results obtained here for the optimal homotopy filter are encouraging, general application of the idea seems problematic. In bearings-only tracking the optimal homotopy can be performed with a vector of five parameters. Even with this relatively small parameter vector the optimal homotopy is difficult to implement, mainly because of the need to evaluate the integrals which define the flow of the parameter values. This issue would be exacerbated in problems where approximation of the posterior requires a larger number of parameters.

6. REFERENCES

- [1] A.H. Jazwinski, *Stochastic Processes and Filtering Theory*, Academic Press, 1970.
- [2] S. Julier, J. Uhlmann, and H.F. Durrant-Whyte, "A new method for the nonlinear transformation of means and covariances in filters and estimators," *IEEE Transactions on Automatic Control*, vol. 45, no. 3, pp. 477–482, 2000.
- [3] A. Doucet, S. Godsill, and C. Andrieu, "On sequential Monte Carlo sampling methods for Bayesian filtering," *Statistics and Computing*, vol. 10, no. 3, pp. 197–208, 2000.
- [4] U.D. Hanebeck, K. Briechle, and A. Rauh, "Progressive Bayes: A new framework for nonlinear state estimation," in *Proceedings of SPIE*, Orlando, USA, 2003, vol. 5099.
- [5] J. Hagmar, M. Jirstrand, L. Svensson, and M. Morelande, "Optimal parameterization of posterior densities using homotopy," in *Proceedings of the International Conference of Information Fusion*, Chicago, USA, 2011.
- [6] F.E. Daum and J. Huang, "Particle flow for nonlinear filters with log-homotopy," in *Proceedings of SPIE*, San Diego, USA, 2008, vol. 6969.
- [7] C. Musso, N. Oudjane, and F. Le Gland, "Improving regularised particle filters," in *Sequential Monte Carlo Methods in Practice*, A. Doucet, N. de Freitas, and N. Gordon, Eds. Springer-Verlag, New York, 2001.
- [8] T. Brehard and J-P. Le Cadre, "Closed-form posterior Cramér-Rao bounds for bearings only tracking," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 42, no. 4, pp. 1198–1223, 2006.