Introduction To Particle Filters For Tracking Applications In The Passive Acoustic Monitoring Of Cetaceans

P. R. White and M. L. Hadley
Signal Processing and Control Group, Institute of Sound and Vibration Research (ISVR),
Highfield, Southampton SO17 1BJ, UK. Emails: (prw,mh1)@isvr.soton.ac.uk.

ABSTRACT

The application of particle filters to two tracking problems in passive acoustic monitoring are discussed. Specifically we describe algorithms for extracting the contours of delphinid whistles and the localization of vocalizing animals in three dimensions using a distributed sensor array. The work is focused on highlighting the potential of particle filters in the analysis of bioacoustic signals. The discussion is based on one particular form of particle filter: the sequential importance resampling filter.

SOMMAIRE

Cette étude porte sur l'application des filtres particulaires à deux problèmes d´extraction d´information en acoustique passive. La description concerne plus spécifiquement deux algorythmes ayant pour objectif l´extraction de contour des sifflements de dauphins et la localisation en trois dimensions d´animaux vocalisant à partir d´un jeu de capteurs répartis localement. L´objectif de ce travail est de mettre en lumière le potentiel des filtres particulaires pour l´analyse de signaux bioacoustiques. Parmi les filtres particulaires, l´accent est mis dans cette étude sur forme particulière de filtre: le filtre à rééchantillonnage d´importance séquentiel.

1. INTRODUCTION

Real-time Passive Acoustic Monitoring (PAM) systems for cetaceans require the integration of many elements. Several of these elements can be cast as tracking problems. In particular this paper considers two such aspects: extracting whistle contours and the estimation of source location using a sensor array. The objective of this work is to highlight the potential of particle filters within the application area as a real-time tracking solution, so the paper is framed in a somewhat pedagogical manner. We avoid details of the theoretical principles under-pinning particle filters, rather we aim to convey the fundamental steps common to particle filters.

The definition of a tracking problem is simply a parameter estimation problem in which the parameter estimates are continually updated; such tasks are also formerly referred to as sequential estimation problems. They have been widely studied in a large range of application areas, including sonar, radar and biomedicine. The classical tool for performing sequential estimation is the Kalman filter and its variants. These methods have been widely, and often successfully, exploited. However their applicability is limited by the underlying assumptions they require.

2. BACKGROUND

The general framework for sequential estimation problems can be expressed as follows. The true value of the parameter vector to be tracked, at time step n, is denoted $\boldsymbol{\theta}_n$. The evolution of this parameter vector is described through a system function, F, such that

$$\boldsymbol{\theta}_n = F\left(\boldsymbol{\theta}_{n-1}, \mathbf{w}_n\right) \tag{1}$$

where $\mathbf{w}_n$ is a vector of random variables specifying the random component of the parameter evolution, it is referred to as either the process or the system noise. Similarly the function G defines the measurement process, where $\mathbf{x}_n$ contains the measured data

$$\mathbf{x}_n = G\left(\boldsymbol{\theta}_n, \mathbf{v}_n\right) \tag{2}$$

in which $\mathbf{v}_n$ represents the measurement noise process. In the general case the functions F and G are non-linear, the noise processes $\mathbf{w}_n$ and $\mathbf{v}_n$ are not necessarily additive and are not distributed according to a Gaussian distribution. Our goal is to estimate the parameter vector $\boldsymbol{\theta}_n$ on the basis of the set of measurements $\mathbf{x}_k$, $k=0,1,\ldots,n$. In order to avoid increasing memory requirements and computational load as n increases, it is natural to seek a recursive solution. That is to say we seek a solution in which the parameter estimate at time n is derived only from knowledge of the parameter estimate at the preceding time step, n-1, and the current measurement $\mathbf{x}_n$. It should be noted that when $k=0$, $\mathbf{x}_0$ is the only information available. This is typically provided by a suitable detection algorithm.

2.1 The Kalman Filter

The Kalman filter is a recursive algorithm which is optimal under simplifying assumptions on the system and measurement models (Arulampalam *et al.*, 2002). Specifically (1) and (2) are simplified so that the system and measurement models are linear and the noise processes are additive and Gaussian. Leading to a model of the form

$$\begin{aligned} \boldsymbol{\theta}_n &= \mathbf{A}_n \boldsymbol{\theta}_{n-1} + \mathbf{w}_n \\ \mathbf{x}_n &= \mathbf{B}_n \boldsymbol{\theta}_n + \mathbf{v}_n \end{aligned} \tag{3}$$

in which $\mathbf{A}_n$ and $\mathbf{B}_n$ are the system and measurement matrices, note that whilst in (3) temporal dependence of these matrices has been assumed, in many applications they are constant. The update equations for the classic Kalman filter are (Bozic, 1979;Zarchan & Musoff, 2005)

$$\begin{aligned} \mathbf{T}_n &= \mathbf{A}_n \mathbf{P}_{n-1} \mathbf{A}_n^t + \mathbf{Q}_n \\ \mathbf{K}_n &= \mathbf{T}_n \mathbf{B}_n^t \left(\mathbf{B}_n \mathbf{T}_n \mathbf{B}_n^t + \mathbf{R}_n \right)^{-1} \\ \mathbf{P}_n &= \mathbf{T}_n - \mathbf{K}_n \mathbf{B}_n \mathbf{T}_n \\ \hat{\boldsymbol{\theta}}_n &= \mathbf{A}_n \hat{\boldsymbol{\theta}}_{n-1} + \mathbf{K}_n \left(\mathbf{x}_n - \mathbf{B}_n \mathbf{A}_n \hat{\boldsymbol{\theta}}_{n-1} \right) \end{aligned} \tag{4}$$

in which $\mathbf{T}_n$ is a temporary matrix (but can be regarded as an a priori estimate of $\mathbf{P}_n$) used to ease the computational load, $\mathbf{Q}_n$ and $\mathbf{R}_n$ are the covariance matrices for the process and measurement noises respectively, $\mathbf{K}_n$ is the Kalman gain matrix, $\mathbf{P}_n$ is the error covariance matrix and $\hat{\boldsymbol{\theta}}_n$ is the vector of parameter estimates at time n.

The Kalman filter is a highly flexible and computationally efficient scheme. But its application is limited to cases where (3) can be regarded as suitable approximation of (1) and (2). Variants on the Kalman filter have been proposed which extend its range of applicability, common examples of these are the extended Kalman filter (EKF) (Zarchan & Musoff, 2005) and unscented Kalman filter (UKF) (Wan & van der Merwe, 2000).

2.2 The Particle Filter

Particle filters provide a general solution to tracking problems of the form described by (1) and (2), without the need to invoke the inherent assumptions associated with the Kalman filter. There are a wide variety of versions of particle filters that have been be defined (Arulampalam *et al.*, 2002;Ristic *et al.*, 2004;Doucet *et al.*, 2001). However the objective of this work is to communicate the opportunities afforded by the use of particle filters in PAM systems, rather than a review of particle filters per se. So we shall concentrate on a simple form of particle filter, specifically we shall discuss Sequential Importance Resampling (SIR) filters. These do not represent the state-of-the-art particle filtering algorithms, but the do provide a good basis for the introduction of the concepts of particle filtering and offer good performance in the examples presented herein.

Consistent with the review character of this publication we provide a mechanistic description of the SIR filter and choose to omit the under-pinning principles, these principles are widely available elsewhere, e.g. (Arulampalam *et al.*, 2002;Ristic *et al.*, 2004;Doucet *et al.*, 2001). The objective here is to provide some insight into how to construct a particle filter and to highlight the flexibility and power that they provide.

Particle filters are also referred to as sequential Monte Carlo algorithms (Doucet *et al.*, 2001) and, as is characteristic of Monte Carlo schemes, they exploit samples drawn from the underlying distributions. Given a set of M parameter estimates at time n-1 which are denoted $\hat{\boldsymbol{\Theta}}_{n-1} = \left\{ \hat{\boldsymbol{\theta}}_{n-1,k} \right\}_{k=1,\ldots,M}$, the basic steps involved in the SIR particle filter are:

i. Update each of the estimates using $\tilde{\boldsymbol{\theta}}_{n,k} = F\left(\hat{\boldsymbol{\theta}}_{n-1,k}, \mathbf{w}_{n,k} \right)$ where $\mathbf{w}_{n,k}$ is a sample from the process noise distribution.

ii. Use the measured data to score each of the new estimates $\tilde{\boldsymbol{\theta}}_{n,k}$ using the likelihood computed via (2) and normalize these scores so that they sum to unity.

iii. Create $\hat{\boldsymbol{\Theta}}_n$ by drawing M samples, with replacement, from $\tilde{\boldsymbol{\Theta}}_n$ according to the scores allocated in step ii.

iv. From the samples $\tilde{\boldsymbol{\Theta}}_n$ form an estimate of the parameter vector.

In the first step the existing estimates are perturbed, in a manner which mimics the effect of process noise, so producing a set of candidate parameter estimates. The particle filter then considers these estimates and scores them according to how well they predict the sample that has just been measured, $\mathbf{x}_n$. This may be explained in the specific case of an additive noise measurement noise model, i.e. in the case where (2) can be expressed in the form:

$$\mathbf{x}_n = G(\boldsymbol{\theta}_n) + \mathbf{v}_n \tag{5}$$

In such cases the scoring is realized by evaluating the likelihood $p_v\left(\mathbf{x}_n - G\left(\tilde{\boldsymbol{\theta}}_{n,k}\right)\right)$ in which p_v is the probability density function of the measurement noise process $\mathbf{v}$. Consequently parameter estimates close to the true value should yield values of $G\left(\tilde{\boldsymbol{\theta}}_{n,k}\right)$ which are close to the measured data, so the have relatively large likelihood. Whereas estimates significantly different form the true value, will (probably) yield measurement estimates very different from the measured value, so yield a low likelihood. The scores are derived from the likelihood by scaling them so that they sum to unity.

Step iii, is realized by selecting the estimates using random sampling according to the estimate's scores. Uniform random variables are used and the probability of selecting a particular estimate is given by its score. The sampling is implemented with replacement, so that estimates with high scores are typically replicated many times. The new samples constitute the set of parameter estimates for starting the next iteration.

The final step is to construct the final parameter estimate. This can be done using one of several principles including: MAP (maximum a priori probability) and minimum mean squared error (MMSE).

3. WHISTLE CONTOUR EXTRACTION

One way in which species classification for delphinids can be achieved is through analysis of their whistles (Oswald *et al.*, 2007). Specifically the contours of the whistles in the time-frequency domain are used as the key features and these contours need to be estimated (extracted) in order to successfully realize such a system. The extraction of such whistles is normally achieved through use of the spectrogram (Oswald *et al.*, 2007;Datta & Sturtivant, 2002;Leprettre & Martin, 2002) although alternative approaches can prove successful (Johansson & White, 2004). The extraction process can be hindered by the presence of overlapping whistles and echolocation clicks as well as potentially low Signal to Noise Ratios (SNRs).

In this work we demonstrate how particle filters can be used as one way to automate this contour extraction process. Other workers have considered applying particle filters to similar problems based on the spectrogram (Dubois *et al.*, 2005;Nagappa & Hopgood, 2006). The work described here takes advantage of the Short Time Fractional Fourier Transform (STFrFT). The STFrFT for the analysis of whistle has been considered elsewhere (Capus & Brown, 2003). It is worth noting that the method we adopt we refer to as a STFrFT, largely in deference to previous work in this application, but it should be noted that the method could be regarded under a number of other signal processing paradigms, most obviously it also exploits the principles behind adaptive basis decomposition methods (Mallat & Zhang, 1993).

Our tracking scheme is based on detecting the maxima of the STFrFT. Consider the k^{th} windowed data segment of the incoming data stream, $x(n)$, denoted $\mathbf{x}_k$ and defined as

$$\mathbf{x}_k = \left[x(kP), x(kP+1), \ldots, x(kP+L-1)\right]^t \tag{6}$$

where P is the number of samples by which the window is shifted between successive analysis windows and L is the window length. The elements of $\mathbf{x}_k$ are denoted $x_k(m)$, m=0,…,L-1. The STFrFT is defined as

$$S(k,\alpha,f) = \left|\sum_{m=0}^{L-1} x_k(m) e^{-2\pi i(f+\alpha t_m/2)t_m}\right|^2 \tag{7}$$

where f denotes centre frequency [Hz], α is the frequency sweep rate [Hz/s] and the local time index, t_m, is defined through $t_m = \left(m - \frac{L}{2}\right)/f_s$, in which f_s denotes the sampling frequency. The STFrFT can be loosely regarded as representing the energy in a signal at a particular time and at a frequency associated with a particular sweep rate.

Evidently the classical short time Fourier transform (the spectrogram) is a special case of (7) in which $\alpha \equiv 0$. The flexibility offered by the STFrFT allows the processing scheme to more accurately model the underlying process. By accommodating linear sweeps the STFrFT can increase the SNR of received signal, assuming that a sweep rate, α, is chosen that is close to that in the received data. The rapid sweep rates that can be observed in odontocete whistles make the use of the STFrFT an attractive option. The use of the STFrFT intrinsically provides estimates of the sweep rate for each analysis window; this additional information can be used to improve tracking performance.

3.1 Particle Filter for Whistle Contour Extraction

The use of particle filters to extract whistle contours requires one to define the system and measurement functions, i.e. (1) and (2). The parameter vector we seek to estimate contains both the frequency and sweep rate and is defined as $\boldsymbol{\theta}=[f, \alpha]^t$. The system model we employ is:

$$\boldsymbol{\theta}_n = \begin{bmatrix} 1 & \frac{P}{f_s} \\ 0 & 1 \end{bmatrix} \boldsymbol{\theta}_{n-1} + \mathbf{w}_n = \mathbf{A}\boldsymbol{\theta}_{n-1} + \mathbf{w}_n \qquad (8)$$

$\mathbf{w}_n$ is a zero mean Gaussian noise with a diagonal covariance matrix, so that system noise on the frequency and sweep rate are uncorrelated with difference variances. This is a standard linear model of the form of (3).

There are several candidate measurements one can use for this system. The one adopted herein, based on the STFrFT, is

$$X_n = \left| \sum_{m=0}^{L-1} x_k(m) e^{-2\pi i(\theta_n(1)+\theta_n(2)t_m/2)t_m} \right|^2 \qquad (9)$$

Note the distinction between the data, $x_k(m)$, and the measurement associated with a parameter X_n. The measurement X_n is a scalar value.

The processing scheme adopted here applies a robust pre-whitening step (Leung & White, 1998) to the incoming data stream, to create $x(n)$; this ensures that the background noise has an approximately flat spectrum of a known level. This pre-whitening allows one to use the value X_n as a proxy for the (unscaled) probability $p(\boldsymbol{\theta}_n \mid x_k)$: large values of X_n relate to highly probable events, whereas small values of X_n relate to events of low probability. This argument is a simplification of the principles lucidly described in detail in (Brethorst, 1988). The non-linear character of (9) favors the use of a particle filter solution.

3.2 Results for Whistle Contour Extraction

This method is applied to a short (1.9 s) section of a whistle recorded from *Tursiops truncatus*. This recording contains features common to many similar recordings. There are trains of echolocation clicks, e.g. between 0.2 and 0.4 s, there are other whistles of varying strength and the primary whistle varies rapidly in frequency and level throughout the recording. The results are shown in Figure 1. The upper frame of this figure shows the original spectrogram, whilst the lower frame depicts the same spectrogram overlaid with a white line showing the estimated frequency contour.

The algorithm has successfully tracked the whistle. There is an initial period, before 0.1 s, where the algorithm provides an estimate which deviates somewhat from the visual track of the whistle. The signal is weak here, but the primary cause for this behavior is the fact that the algorithms require some time to initialize, to “burn in”. The estimated track also deviates at around 0.6 s when the whistle’s amplitude temporarily reduces significantly. Accepting these minor deviations it is encouraging to note that the algorithm has successfully tracked the whistle even during the rapid frequency jump occurring shortly after 0.2 s and ending shortly before 0.4 s. This is despite the signal being partly obscured by a click train. This is particularly gratifying since such jumps are characteristic of *T. truncatus* and one would seek to avoid classifying such a jump as two separate whistles.

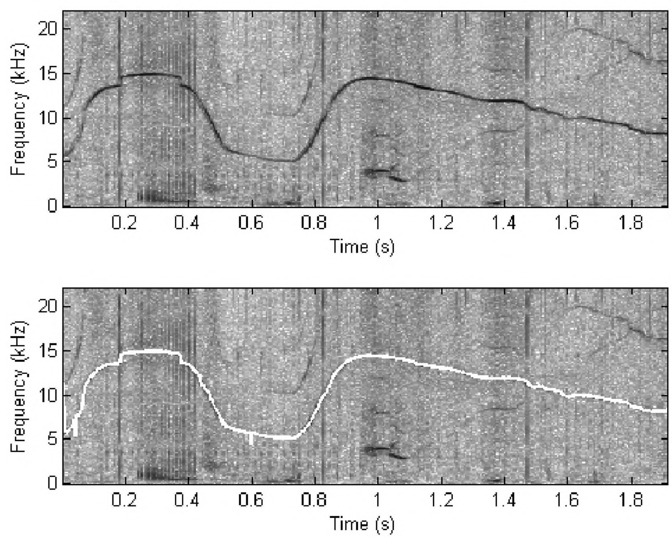

Figure 1: Results of contour extraction for a *Tursiops truncatus* whistle based on a particle filter. Upper frame shows the spectrogram, the lower frame shows the same plot with a white line overlaid to show the contour estimate.

4. SOURCE LOCALISATION

The second problem this paper considers is that of tracking the location of a vocalizing animal in three dimensions using a hydrophone array. The most suitable signals for performing such localizations are echolocation clicks, but other vocalizations can be used. In this example we use echolocation clicks from a sperm whale, *Physeter macrocephalus*.

The localization problem can be described as follows. A set of acoustic sensors (in this case hydrophones) are located at known positions, $\mathbf{r}_m$, in an environment. The source, is at an unknown location $\mathbf{s}$, and emits a sound (a vocalization) which propagates through a known medium. The phrase "known medium" is intended to highlight the assumption that the propagation time from one point in the medium to a second point can computed, implying knowledge of (at least) the sound speed profile. This model we shall denote $\mathcal{M}$. Using the received signals one is able to compute the delays, $\tau(m,n)$, observed between the vocalizations being detected on the pair of hydrophones m and n. Collecting all of these delays into a vector $\boldsymbol{\tau}$ allows one to express the problem as: given the measured delays $\boldsymbol{\tau}$, a model of the acoustic environment, $\mathcal{M}$, and the sensor locations, $\mathbf{r}$, can one infer the source location?

The choice of the acoustic model $\mathcal{M}$ is an important factor controlling the accuracy of the resulting estimated source locations. The use of a simple model should result in an efficient algorithm, but the estimates may be subject to considerable error. The use of models that accurately capture the propagation of sound in the ocean is clearly advisable, but their use is often limited by the absence of complete knowledge of the physical parameters required to specify such a model. The form of the model used does not directly impact the following discourse. This problem has been treated by a large number of authors as a non-sequential estimation problem, e.g. (Spiesberger, 2001;Thode, 2004;White *et al.*, 2006). In the following we translate the problem into a tracking, sequential estimation, task and present a particle filter based solution.

4.1 Particle Filter for Source Localization

The underlying model when using a particle filter for localization is relatively straightforward. The unknown parameter vector, $\boldsymbol{\theta}$, contains the source co-ordinates, for example expressed in Cartesian co-ordinates. The system matrix aims to model how the animal moves through the medium. Various methods can be used to impose models that are appropriate for the known behavioral parameters. For example one can seek to impose maximum swim rates, or rates of ascent and descent. With the objective of retaining simplicity we stick with a simple random walk model:

$$\boldsymbol{\theta}_p = \boldsymbol{\theta}_{p-1} + \mathbf{w}_p \quad (10)$$

where $\mathbf{w}_p$ represents a vector of independent, zero mean, Gaussian white noise. The underlying assumptions in (10) are very limited, it assumes that the current location is the just a random perturbation from the preceding location; Further note that the vocalizations typically occur at irregular intervals. The subscript p denotes data associated with the p^{th} such vocalization. A shortcoming of the random walk model described by (10) is that it does not account for this irregular sampling. It is reasonable to increase the standard deviation of $\mathbf{w}_p$ in proportion to the interval between irregular vocalizations, reflecting the fact that an animal is likely to have moved a greater distance in longer intervals than short ones.

The measurement model is simply:

$$\boldsymbol{\tau}_p = \mathcal{M}\left(\boldsymbol{\theta}_p, \mathbf{r}\right) + \mathbf{v}_p \quad (11)$$

The function $\mathcal{M}$ is typically a highly non-linear function. It is this that again makes the particle filter an attractive processing option. The measurement noise $\mathbf{v}_n$ is modeled using a long-tailed distribution, such as a Laplacian distribution. The advantage of this is that it models the occasional failure of the delay estimation operation. The presence of strong reflectors can lead to some isolated delay estimates with large errors. By employing a measurement noise model with a long-tailed distribution such outliers are penalized less than would be the case if a Gaussian model was used for the noise.

4.2 Results for Source Localization

The algorithm outline in the preceding subsection has been applied to data obtained from an echo-locating sperm whale using bottom mounted hydrophones. Specifically, the data used in this study was the second data set supplied for the 2005 Workshop on Detection and Localization of Marine Mammals using Passive Acoustics held in Monaco (Adam, 2006).

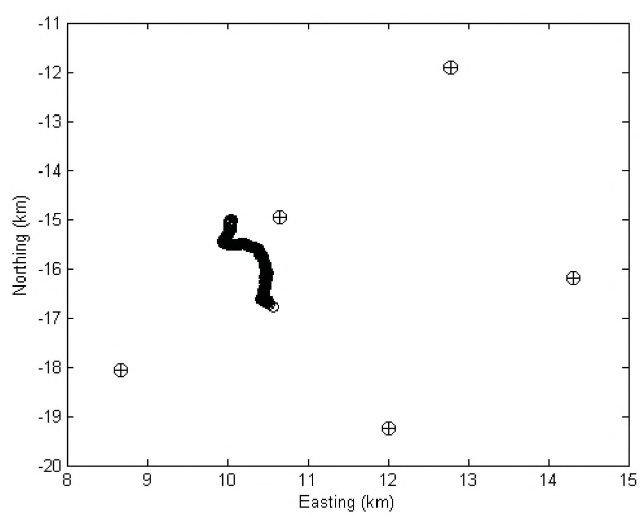

Figure 2: Estimated source locations using particle filter. Open circles 'O' indicate the individual estimates. Crossed circles '⊕' indicate the sensor locations

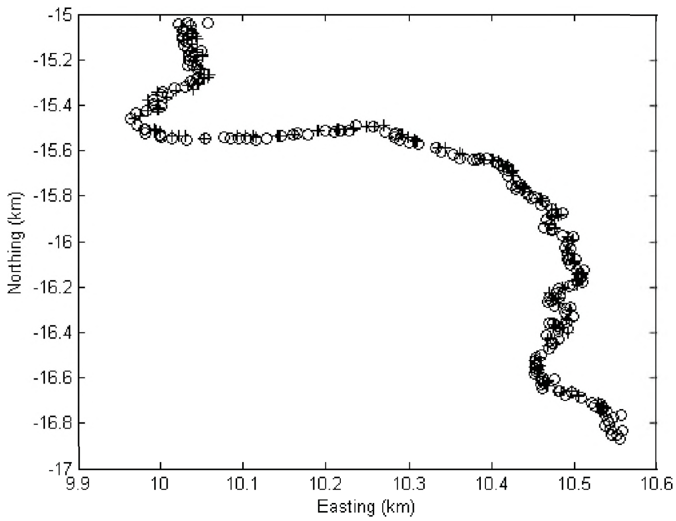

Figure 3: Comparison of estimated source locations using particle filter and non-sequential estimation.
Open circles 'O' indicate the results of particle filtering.
Crosses '+' indicate estimates from non-sequential estimator

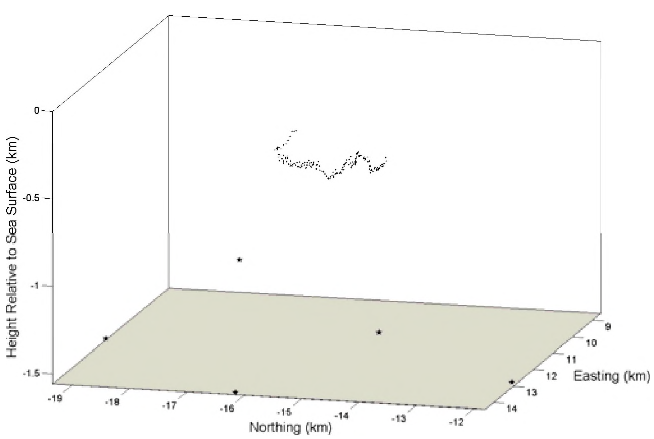

Figure 4: Estimated source location in three dimensions.
Sensor location locations are indicated by the black stars (⋆)

The results depicted here are based on the same delay estimates as those employed in (White *et al.*, 2006). The results are obtained assuming a linear sound speed profile. Figure 2 depicts the result of applying the particle filter to this data set. Whilst Figure 3 shows the same results on an expanded scale, also shown in this plot are the results of the algorithm presented in (White *et al.*, 2006) obtained on the same data set using a propagation model with a constant sound speed. Comparing the results from the particle filter with those from the estimation method, in (White *et al.*, 2006), demonstrates the potential of the particle filter as a real-time tracking solution of comparable performance given the same time delay estimates.

The results shown in Figures 2 and 3 are consistent with those obtained in (Adam, 2006) and from Figure 3 we see that the particle filter results and those from the non-sequential scheme are very close to each other. In Figure 4 these results are show in three dimensions.

5. DISCUSSION

Particle filters provide a general, powerful and flexible tool for solving tracking problems. The results herein demonstrate that the solutions achieved are of high quality, despite the rather crude nature of the particle filter algorithms used and the fact that we have, in general, avoided including all of the available prior information in the interests of retaining simplicity.

The good performance of particle filters is commonly realized at the cost of a large computational burden being incurred. A key parameter in controlling this cost is the number of particles M employed. The larger the number of particles, the better the solution but the greater the computational burden imposed. A second key parameter that affects performance, but is not related to computational burden, is suitability of the choice of the distribution width. This should be chosen to be representative of the change expected to occur in the state vector between measurements. In the source localization application this would be relative to the typical swim speed of the species of animal to be tracked.

The initial estimate of the state vector is derived depending on the application. For the whistle contour extraction once the presence of a whistle is detected each STFrFT bin is weighted according to a uniform distribution. Here this is possible because the states are discrete and therefore the number is relatively limited. In the localization application the number of possible initial states is much greater, therefore a cost function minimization estimation scheme was utilized to provide the first estimate.

The implementations employed herein both used 1000 particles. In the case of the localization the computational cost that implied was consistent with a real-time implementation, even with the algorithm implemented in MATLAB®. This is in part because the filter only needs updating approximately once per second. The computational load would also escalate significantly if a more detailed propagation model is used.

The real-time implementation of the contour extraction algorithm requires modification of the algorithm presented here as in its current form it is probably too computationally demanding for simple real-time implementation. One could exploit the potential for parallel implementation inherent in particle filters, but this dramatically increases the issues ssociated with implementation.

REFERENCES

1. Adam,O. (2006) Special issue: Detection and localization of marine mammals using passive acoustics. *Applied Acoustics,* **67**, 1057-1058.

2. Arulampalam,M.S., Maskell,S., Gordon,N. & Clapp,T. (2002) A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing,* **50**, 174-188.

3. Bozic,S.M. (1979) *Digital and Kalman Filtering*. Edward Arnold, London.

4. Brethorst,G.L. (1988) *Bayesian spectrum analysis and parameter estimation*. Springer-Verlag, Berlin.

5. Capus,C. & Brown,K. (2003) Short-time fractional Fourier methods for the time-frequency representation of chirp signals. *Journal of the Acoustical Society of America,* **113**, 3253-3263.

6. Datta,S. & Sturtivant,C. (2002) Dolphin whistle classification for determining group identities. *Signal Processing,* **82**, 251-258.

7. Doucet,A., de Freitas,N. & Gordon,N. (2001) *Sequential Monte Carlo Methods in Practice*. Springer Verlag.

8. Dubois,C., Davy,M. & Idier,J. (2005) Tracking of time-frequency components using particle filtering. In: *International Conference on Speech and Signal Processing 2005,* p. IV-9-IV-12.

9. Johansson,A.T. & White,P.R. (2004) Detection and characterization of marine mammal calls by parametric modelling. *Canadian Acoustics,* **32**, 83-92.

10. Leprettre,B. & Martin,N. (2002) Extraction of pertinent subsets from time-frequency representations for detection and recognition purposes. *Signal Processing,* **82**, 229-238.

11. Leung,T.S. & White,P.R. (1998) Robust estimation of oceanic background spectrum. In: *Mathematics in signal processing IV,* McWhirty J.G. & Proudler I.K. (eds.), pp. 369-382. Clarendon Press, Oxford.

12. Mallat,S.G. & Zhang,Z.F. (1993) Matching Pursuits with Time-Frequency Dictionaries. *Ieee Transactions on Signal Processing,* **41**, 3397-3415.

13. Nagappa,S. & Hopgood,J.R. (2006) Frequency Tracking of Biological Waveforms. In: *IMA Seventh International Conference On Mathematics In Signal Processing,* p. 12.

14. Oswald,J.N., Rankin,S., Barlow,J. & Lammers,M.O. (2007) A tool for real-time acoustic species identification of delphinid whistles. *Journal of the Acoustical Society of America,* **122**, 587-595.

15. Ristic,B., Arulampalam,M.S. & Gordon,N. (2004) *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Artech House.

16. Spiesberger,J.L. (2001) Hyperbolic location errors due to insufficient numbers of receivers. *Journal of the Acoustical Society of America,* **109**, 3076-3079.

17. Thode,A. (2004) Tracking sperm whale (Physeter macrocephalus) dive profiles using a towed passive acoustic array. *Journal of the Acoustical Society of America,* **116**, 245-253.

18. Wan,E.A. & van der Merwe,R. (2000) The unscented Kalman filter for nonlinear estimation. In: *IEEE Symposium on Adaptive Systems for Signal Processing, Communications and Control,* pp. 153-158.

19. White,P.R., Leighton,T.G., Finfer,D.C., Powles,C. & Baumann,O.N. (2006) Localisation of sperm whales using bottom-mounted sensors. *Applied Acoustics,* **67**, 1074-1090.

20. Zarchan,P. & Musoff,H. (2005) *Fundamentals of Kalman filtering: a practical approach*, 2nd edn. AIAA.