

FREQUENCY TRACKING USING MONTE CARLO METHODS: APPLICATION TO BAT ECHOLOCATION SIGNALS

Sharad Nagappa and James Hopgood

Institute for Digital Communications, University of Edinburgh,
Edinburgh - EH9 3JL, United Kingdom
email: S.Nagappa@ed.ac.uk, James.Hopgood@ed.ac.uk

ABSTRACT

This paper presents a sequential Monte Carlo method for tracking an unknown varying number of time-varying frequencies. A reversible-jump sampler is used to implement model-order determination. It is shown that for a linear-in-the-amplitudes observation model in white Gaussian noise, the amplitudes and noise variance can be analytically marginalised out of the posterior distribution resulting in a reduced dimension state estimation problem. A sum of linear chirps model is chosen as a local observation model and this basis is used to determine the instantaneous frequency of the signal. We present frequency tracking results from synthetic as well as field-recorded bat echolocation calls and compare results with Fourier based frequency tracking.

1. INTRODUCTION

Time frequency analysis of biosonar signals can be performed using techniques based around the short-time Fourier transform (STFT) or wavelets, with further analysis relying on this time-frequency representation. This work addresses the issue of instantaneous frequency estimation by direct inference from the raw data and without any preprocessing. A sequential Bayesian approach is adopted to solve this problem. The use of a Bayesian approach facilitates incorporation of prior information into the estimator, allowing the estimator to use all the information available.

Figure 1 [1] illustrates the time-varying signal structure as a bat attempts to first locate and identify potential targets (usually insects, as a source of food), then tracks the target until it is finally captured. In the figure, there are certain obvious changes in the signal structure over the *search-approach-terminal* phases: the duration of each call decreases, multiple harmonics may be introduced, and the pulse repetition rate increases. There is very little quantitative analysis of the bat calls in the literature that exists beyond such a descriptive analysis. The aim of this work is to then provide a robust method to determine and track the time-varying frequency of multiple components of such a signal, thus providing a basis for quantitative analysis of these signals. Such an analysis would be useful to biologists studying echolocation in bats.

The implementation of a sequential approach offers several advantages over a batch approach. A batch approach necessitates a model for the time-frequency structure of the bat call which varies across species. In addition, each call can have different amplitude envelopes, and multiple harmonics, which may not last the entire duration of the signal. The batch approach can result in a difficult estimation problem with many parameters.

To overcome the problem of “parameter-bloat”, we choose to sequentially estimate the frequencies of the subcomponents present in the signal. As a result, no specific model needs to be considered for different species while estimating the frequency. Using these

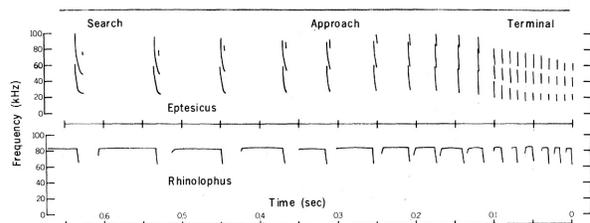


Figure 1: Diversity of echolocation call structure

frequency estimates, model determination and parameter estimation can be performed subsequently.

One approach for estimating the instantaneous frequency content of a signal would be to locate peaks in the spectrogram of a signal. Such a simplistic approach is not necessarily either reliable or robust. Approaches for time-varying frequency tracking have been proposed previously [2, 3, 4] taking into account smooth changes in the frequency trajectory. STFT-based time-frequency tracking [2] is implemented by applying a particle filter [5] to track the peaks present in the spectrogram of a signal. The method looks at the Fourier transform over short blocks of the observation sequence and estimates in each block the number of components present, as well as their frequencies, amplitudes, and noise variance.

The method described here is related to the harmonic tracking algorithm described in [4]. In their work, a jump Markov system (JMS) is used to detect multiple harmonic components. A Rao-Blackwellised particle filter (RBPF) [5] is used to integrate out amplitude parameters in their model. In contrast, we investigate an alternative approach where the amplitude and noise variance parameters are analytically marginalised out due to the use of a linear-in-the-amplitudes observation model in Gaussian noise [6, 7]. Further, reflecting the data set of interest, viz., bat echolocation calls, we adopt a linear chirp as opposed to a sinusoidal basis.

When dealing with biosonar signals, the number of components present in the signal is a time-varying parameter. Reversible-jump Markov chain Monte Carlo (RJMC) methods [8] have previously been used to estimate an unknown number of frequencies for the stationary-frequency case [9]. Within a sequential Monte Carlo (SMC) framework, we demonstrate that the idea can be extended and applied to the problem of frequency-tracking when the number of frequency components is time-varying. While a JMS can be applied to detect a varying number of components, the RJMC sampler is incorporated into a SMC framework since it improves filter performance by inhibiting unlikely moves [10].

2. SIGNAL MODEL

The rate of change of phase of a signal provides the instantaneous frequency of the signal. For a multicomponent signal, it is not the overall rate of change of phase that we are interested in, but the combination of frequencies which are present at that instant. We redefine “instantaneous frequency” here so that each subcomponent of the signal has its own instantaneous frequency.

In order to estimate the instantaneous frequencies, we slide a window over the observation sequence and estimate the frequency

IDCOM is in the Joint Research Institute for Signal and Image Processing, a member of the Edinburgh Research Partnership in Engineering and Mathematics. This work is sponsored by the EPSRC under their basic technology grant, “Biologically Inspired Acoustic Systems”. Bat recordings used in this work were obtained from Dr Dean Waters, University of Leeds.

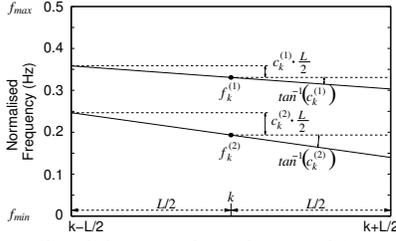


Figure 2: Over the sliding window, the signal is modelled as the sum of linear chirps. In this case, there are two linear chirps with centre frequencies $(f_k^{(1)}, f_k^{(2)})$, and chirp-rates $(c_k^{(1)}, c_k^{(2)})$.

components present in the windowed section of the observations. The time varying nature of each frequency component is modelled along the lines of the equation

$$f_{k+1} = g(f_k) + w_k \quad (1)$$

where f_k is the frequency at time k , $g(\cdot)$ is a function to update the frequency from the previous time instant, and $w_k \sim \mathcal{N}(0, \sigma_w^2)$.

Within the sliding window, the observation segment is modelled locally as the sum of P_k linear chirps with frequencies $\mathbf{f}_k = [f_k^{(1)}, f_k^{(2)}, \dots, f_k^{(P_k)}]^T$ and chirp-rates $\mathbf{c}_k = [c_k^{(1)}, c_k^{(2)}, \dots, c_k^{(P_k)}]^T$. Let L be the length of the sliding window defined such that k is at the centre of the window. Over this window, the observation x_t at time $t \in \{k - L/2, \dots, k + L/2\}$, assuming L odd, is

$$x_t = \sum_{p=1}^{P_k} \left(a_k^{(p)} \cos \phi_{k,t}^{(p)} + b_k^{(p)} \sin \phi_{k,t}^{(p)} \right) + n_t \quad (2)$$

$$\phi_{k,t}^{(p)} = 2\pi \left(f_k^{(p)} - c_k^{(p)} \frac{L}{2} + \frac{c_k^{(p)}}{2} t \right) t \quad (3)$$

where $\mathbf{a}_k = [a_k^{(1)}, a_k^{(2)}, \dots, a_k^{(P_k)}, b_k^{(1)}, b_k^{(2)}, \dots, b_k^{(P_k)}]^T$ are the amplitudes of the cosine and sine amplitudes of the subcomponents, and $n_t \sim \mathcal{N}(0, \sigma_{n,k}^2)$ is zero-mean white Gaussian observation noise with variance $\sigma_{n,k}^2$. Setting $c_k^{(p)}$ to zero in equation 3 reduces the local model to a sum of sinusoids (SoS) model. Figure 2 illustrates the parameters used in the local model.

The selection of a suitable window length depends on the application under consideration. Shorter windows allow us to deal with highly non-stationary frequencies since the variation across a short window will be less. However, the short window also leads to greater variance in the parameter estimate. Long windows, by contrast, yield lower variance estimates, but may yield worse estimates due to frequency non-stationarity within the window. This trade-off must be considered when selecting a window of suitable length.

3. THE LIKELIHOOD FUNCTION

Since the noise term in equation 2 is zero-mean white Gaussian noise, the likelihood function can easily be written down. The likelihood function is evaluated over the windowed section of the observations. Let $\mathbf{x}_k = [x_{k-L/2}, \dots, x_{k+L/2}]^T$ be the vector of observations; $\mathbf{n}_k = [n_{k-L/2}, \dots, n_{k+L/2}]^T$ be the noise.

The model in equation 2 can be written in matrix vector form by defining the following matrices:

$$\Phi_k = \begin{bmatrix} \phi_{k,-L/2}^{(1)} & \phi_{k,-L/2}^{(2)} & \dots & \phi_{k,-L/2}^{(P_k)} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{k,L/2}^{(1)} & \phi_{k,L/2}^{(2)} & \dots & \phi_{k,L/2}^{(P_k)} \end{bmatrix}$$

$$\mathbf{G}_k = [\cos(\Phi_k) \quad \sin(\Phi_k)]$$

where $\phi_{k,t}$ is defined in equation 3; $\cos(\cdot)$ and $\sin(\cdot)$ operate element-wise on the matrix Φ_k to produce an augmented $L \times 2P_k$ \mathbf{G}_k matrix. The signal model can be rewritten in the form of the

general linear model (GLM) [6, 7] as

$$\mathbf{x}_k = \mathbf{G}_k \mathbf{a}_k + \mathbf{n}_k \quad (4)$$

If $\psi_k = \{P_k, \mathbf{f}_k, \mathbf{c}_k, \mathbf{a}_k, \sigma_{n,k}^2\}$, the likelihood function is

$$p(\mathbf{x}_k | \psi_k) = \frac{1}{(2\pi\sigma_{n,k}^2)^{L/2}} \exp \left[-\frac{\|\mathbf{x}_k - \mathbf{G}_k \mathbf{a}_k\|^2}{2\sigma_{n,k}^2} \right] \quad (5)$$

3.1 Parameter Reduction using Marginalisation

From equation 5, it is possible to obtain a likelihood distribution with the amplitude and noise variance terms marginalised over [7]. We would like to remove the dependence of equation 5 on $\Upsilon_k = \{\mathbf{a}_k, \sigma_{n,k}^2\}$, thus reducing the parameter space. The frequency and chirp-rate parameters can then be estimated, for example, using a non-linear search. Since the observations are linear in the amplitude parameters, these can be estimated separately once the frequencies and chirp-rates are determined.

A multivariate normal distribution is defined on the $2P_k$ amplitude parameters, $\mathbf{a}_k \sim \mathcal{N}(\mathbf{0}, \Sigma_k)$, $\Sigma_k = \sigma_{n,k}^2 \delta_k^2 \mathbb{I}_{2P_k}$, where \mathbb{I}_Q is the $Q \times Q$ identity matrix. An extra parameter, δ_k^2 , is introduced, which is indicative of the SNR of the signal. The noise variance is assumed to follow an inverse-gamma distribution, $\sigma_{n,k}^2 \sim \mathcal{IG}(\alpha_n, \beta_n)$, with scale and shape parameters (α_n, β_n) . Thus, more explicitly,

$$p(\mathbf{a}_k | P_k, \sigma_{n,k}^2, \delta_k^2) = \frac{1}{(2\pi\delta_k^2\sigma_{n,k}^2)^{P_k}} \exp \left(-\frac{1}{2} \mathbf{a}_k^T \Sigma_k^{-1} \mathbf{a}_k \right)$$

$$p(\sigma_{n,k}^2 | \alpha_n, \beta_n) = \frac{\alpha_n^{\beta_n}}{\Gamma(\beta_n)} (\sigma_{n,k}^2)^{-(\beta_n+1)} \exp \left(-\frac{\alpha_n}{\sigma_{n,k}^2} \right)$$

The marginalised likelihood function, $p(\mathbf{x}_k | \psi'_k)$, where $\psi'_k = \{P_k, \mathbf{f}_k, \mathbf{c}_k, \delta_k^2\}$, is then obtained from:

$$p(\mathbf{x}_k | \psi'_k) = \int p(\mathbf{x}_k, \Upsilon_k | \psi'_k) d\Upsilon_k = \int p(\mathbf{x}_k | \psi'_k, \Upsilon_k) p(\Upsilon_k) d\Upsilon_k$$

$$\propto \frac{(\mathbf{x}_k^T \mathbf{x}_k + 2\alpha_n - \mathbf{x}_k^T \mathbf{G}_k \mathbf{F}_k^{-1} \mathbf{G}_k^T \mathbf{x}_k)^{-(\frac{L}{2} + \beta_n)}}{\delta_k^{2P_k} \sqrt{\det(\mathbf{F}_k)}} \quad (6)$$

where $\mathbf{F}_k = \mathbf{G}_k^T \mathbf{G}_k + \delta_k^{-2} \mathbb{I}_{2P_k}$.

Priors on the remaining parameters are specified as uniform on the frequency and chirp rate, an inverse-gamma on δ_k^2 , $p(\delta_k^2) = \mathcal{IG}(\alpha_\delta, \beta_\delta)$, and a Poisson distribution on P_k , $p(P_k) = p(P_k | \lambda)$.

3.2 Parameter Estimation

To obtain, for example, the maximum likelihood estimate, a difficult multidimensional search for the function maximum is necessary. This is further complicated by the fact that the number of parameters to be estimated is a parameter that needs to be estimated as well. The complexity can be reduced by assuming that the signal does not change significantly between time k and $k+1$ and exploiting this within a sequential Bayesian framework. We adopt this solution to estimate the signal parameters by modelling the sequential update of the parameters along the lines of equation 1.

Particle filtering methods offer a framework for the implementation of the recursive Bayesian filter, and we use this framework to estimate and track the frequency content of the time-varying signal.

4. PARTICLE FILTERING FRAMEWORK

A sequential importance-sampling resampling (SIR) particle filter [5] is used to perform online frequency-tracking. The particle filter would normally be used to estimate all the state parameters. However, we will demonstrate next that for the GLM, within a sequential framework, the amplitude and noise variance terms may be analytically marginalised out of the posterior distribution.

Since many problems can be placed within the context of the GLM, this marginalisation has wide application. The benefit is that

while the marginalisation introduces an extra parameter, at the same time $2P_k + 1$ parameters are eliminated from the likelihood function.

A particle filter approximates the posterior distribution by a set of weighted samples using sequential importance-sampling (SIS). Let ω_k be the unknown state at time k , $\Omega_k = \{\omega_j\}_{j=0}^k$, and \mathbf{x}_k be the observation, $\mathbf{X}_k = \{\mathbf{x}_j\}_{j=0}^k$. Then $p(\Omega_k | \mathbf{X}_k)$, the posterior distribution, is approximated by a set of N discrete weighted samples, or particles, $\{\Omega_k^{(i)}, w_k^{(i)}\}_{i=1}^N$, where $w_k^{(i)}$ is the weight associated with the sample $\Omega_k^{(i)}$.

$$p(\Omega_k | \mathbf{X}_k) \approx \sum_{i=1}^N w_k^{(i)} \delta(\Omega_k - \Omega_k^{(i)}) \quad (7)$$

The weight $w_k^{(i)}$ evaluates to

$$w_k^{(i)} \propto \frac{p(\mathbf{x}_k | \omega_k^{(i)}) p(\omega_k^{(i)} | \omega_{k-1}^{(i)}) p(\Omega_{k-1}^{(i)} | \mathbf{X}_{k-1})}{q(\omega_k^{(i)} | \Omega_{k-1}^{(i)}, \mathbf{X}_k) q(\Omega_{k-1}^{(i)} | \mathbf{X}_{k-1})} \quad (8)$$

$$= w_{k-1}^{(i)} \frac{p(\mathbf{x}_k | \omega_k^{(i)}) p(\omega_k^{(i)} | \omega_{k-1}^{(i)})}{q(\omega_k^{(i)} | \Omega_{k-1}^{(i)}, \mathbf{X}_k)} \quad (9)$$

where $w_{k-1}^{(i)}$ is the weight at the previous instant $k-1$, the likelihood $p(\mathbf{x}_k | \omega_k^{(i)})$ is obtained from equation 5, $p(\omega_k^{(i)} | \omega_{k-1}^{(i)})$, the transition prior, is obtained according to the state update equations (section 4.1). $q(\omega_k^{(i)} | \Omega_{k-1}^{(i)}, \mathbf{X}_k)$ is the sampling distribution dependent on previous states and observations.

To marginalise the parameters $\Upsilon_k = \{\mathbf{a}_k, \sigma_{n,k}^2\}$ from the posterior, we must integrate over these parameters in equation 7. It is possible to use Rao-Blackwellisation to marginalise the linear amplitude parameters [4], however, we adopt a different approach in analytically marginalising the amplitude and noise variance parameters. Rather than assume a particular model on the Υ_k parameters (e.g. random walk, as in [4]), we assume that the parameters are independent across windows which allows us to easily carry out the marginalisation. The marginalised posterior is written as:

$$p(\Omega'_k | \mathbf{X}_k) \approx \int \dots \int \sum_{i=1}^N w_k^{(i)} \delta(\Omega'_k - \Omega_k^{(i)}) d\Upsilon_{0:k} \quad (10)$$

$$\approx \sum_{i=1}^N w_k^{(i)} \delta(\Omega'_k - \Omega_k^{(i)})$$

where $\Omega'_k = \Omega_k - \{\Upsilon_{0:k}\}$ and $w'_k = \int \dots \int w_k d\Upsilon_{0:k}$.

From equation 10 the posterior distribution can be obtained by the integral of the weight in equation 8 with respect to the nuisance parameters giving rise to a 'marginalised' weight w'_k . The derivation of this marginalised weight update is listed in appendix A.

We utilise the priors listed in section 3.1 and set $\omega'_k = \psi'_k = \{P_k, \mathbf{f}_k, \mathbf{c}_k, \delta_k^2\}$. Under the assumption that $p(\Upsilon_k | \Upsilon_{k-1}) = p(\Upsilon_k)$, the marginalised likelihood $p(\mathbf{x}_k | \omega'_k)$ is identical to equation 6. In addition, the sampling distribution in the SIS step is chosen as the transition prior resulting in the simplified weight update equation

$$w'_k \propto w'_{k-1} p(\mathbf{x}_k | \psi_k^{(i)}) \quad (11)$$

In the approach adopted here, the SIR filter tracks frequencies which have already been detected but does not detect any change in the number of components. A RJMCMC move is subsequently introduced to detect changes in the number of signal components.

4.1 State Update Equations

The state parameters ψ' need to be updated from time t_k at index k to time $t_k + \Delta_t$ at $k+1$.

$$\begin{aligned} \mathbf{f}_{k+1} &= \mathbf{f}_k + \mathbf{c}_k \Delta_t + \mathbf{v}_{f,k} \\ \mathbf{c}_{k+1} &= \mathbf{c}_k + \mathbf{v}_{c,k} \\ \log(\delta_{k+1}^2) &= \log(\delta_k^2) + v_{\delta,k} \end{aligned} \quad (12)$$

where $\mathbf{v}_{f,k} = [v_{f,k}^{(1)}, \dots, v_{f,k}^{(P_k)}]^T$, $\mathbf{v}_{c,k} = [v_{c,k}^{(1)}, \dots, v_{c,k}^{(P_k)}]^T$, $v_{f,k}^{(p)} \sim \mathcal{N}(0, \sigma_f^2)$, is the process noise for update of the p^{th} frequency, $v_{c,k}^{(p)} \sim \mathcal{N}(0, \sigma_c^2)$, is the process noise for update of the p^{th} chirp-rate, $v_{\delta,k} \sim \mathcal{N}(0, \sigma_\delta^2)$, is the process noise for update of the δ_k term. Since the signal is modelled as a linear chirp, $f_k^{(p)}$ can be expected to increase by the amount $\mathbf{c}_k \Delta_t$ at time index $k+1$, and the frequency update is formulated to mirror this.

Changes to the model order P_k are not reflected in the state space equations. The state space equations reflect the update behaviour of the SIR particle filter which is unable to cope with changes in the model order. The model order term is thus updated according to the RJMCMC sampler (section 4.3).

4.2 Instantaneous Parameter Estimates

The particle filter produces an approximation of the multidimensional distribution over the multiple signal parameters. Each particle contains an unordered set of frequency and chirp-rate parameters. As a result, taking a sample mean does not yield a meaningful estimate of the instantaneous frequency content of the signal. We choose the MAP estimate as being representative of the instantaneous signal parameters. The instantaneous estimate is written as

$$\hat{\psi}'_k = \arg \psi'_k \max p(\psi'_k | \mathbf{X}_k) \quad (13)$$

Using the estimated ψ'_k parameters, it is possible to obtain an estimate of the amplitude parameters using a least squares solution:

$$\mathbf{a}_k = (\mathbf{G}_k^T \mathbf{G}_k)^{-1} \mathbf{G}_k^T \mathbf{x}_k$$

4.3 Model-Order Determination using Reversible-Jump Markov Chain Monte Carlo

An alternative to using JMS for model-order selection is to utilise a RJMCMC step in a particle filter which offers certain benefits. The reversible-jump sampler accepts a new state according to an acceptance probability. This accept/reject mechanism ensures that good particles are not lost, thus reducing the variance of the weights and limiting particle degeneracy [10]. This benefit, however, comes at the cost of applying the sampler to each particle in the filter.

In contrast with general MCMC methods, the reversible-jump sampler does not require a burn-in period when used in the SMC context. The reversible-jump sampler requires the burn-in period so as to sample from the limiting distribution. Since the samples from the SIR filter are already distributed accordingly, a single iteration is sufficient for model-order determination [11].

The reversible-jump move updates the number of components tracked by each particle at time k . Moves between different dimension spaces are performed using birth, death, and update moves [8, 9] with respective probabilities $\{b_k, d_k, u_k\}$. Let $\{P_k, \Theta_k\}$ denote the current state, and $\{P'_k, \Theta'_k\}$, the proposed state, where P_k is the number of components and $\Theta_k = \psi'_{k-\{P_k\}}$. The new state will be accepted according to an acceptance ratio r ,

$$r = \underbrace{\frac{p(P'_k, \Theta'_k | \mathbf{x}_k, \Lambda)}{p(P_k, \Theta_k | \mathbf{x}_k, \Lambda)}}_{\text{posterior ratio}} \underbrace{\frac{d(P_k, \Theta_k | P'_k, \Theta'_k)}{d(P'_k, \Theta'_k | P_k, \Theta_k)}}_{\text{proposal ratio}} \underbrace{\mathbf{J}}_{\text{Jacobian}}$$

where $d(\cdot | \cdot)$ denotes the conditional proposal distribution for the parameters. The Jacobian term evaluates to unity for birth and death moves. Once a particular move type is selected, a new state is proposed which is then accepted with an acceptance probability $\alpha = \min\{1, r\}$. Algorithm 1 lists the implementation of the filter.

5. RESULTS

We first present a comparison of the SoS and SoLC models. This will illustrate the gains achieved from using the more complicated model. Since the target class of signals is bat echolocation calls

Algorithm 1 Tracking a multicomponent signal

Initialisation ($k = 0$):

$(P_0^{(i)}, \mathbf{f}_0^{(i)}, \mathbf{c}_0^{(i)}, \delta_k^{2(i)}) \sim p(P, \mathbf{f}, \mathbf{c}, \delta^2) \big|_{i=1}^N$ where i is the particle index, and N is the number of particles used.

1. Increment $k \leftarrow k + 1$

2. Process-update:

$$\begin{aligned} P_k^{(i)} &= P_{k-1} \big|_{i=1}^N \\ \mathbf{f}_k^{(i)} &\sim p(\mathbf{f} | P_{k-1}^{(i)}, \mathbf{f}_{k-1}^{(i)}, \mathbf{c}_{k-1}^{(i)}) \big|_{i=1}^N \\ \mathbf{c}_k^{(i)} &\sim p(\mathbf{c} | P_{k-1}^{(i)}, \mathbf{c}_{k-1}^{(i)}) \big|_{i=1}^N \\ \delta_k^{2(i)} &\sim p(\delta^2 | (\delta_{k-1}^{2(i)}) \big|_{i=1}^N \end{aligned}$$

3. Evaluate importance weight (equation 11):

$$\hat{w}_k^{(i)} \propto w_{k-1}^{(i)} p(\mathbf{x}_k | \psi_k^{(i)})$$

4. Normalise weights:

$$w_k^{(i)} = T^{-1} \hat{w}_k^{(i)} \big|_{i=1}^N; \quad T = \sum_{i=1}^N \hat{w}_k^{(i)}$$

5. Obtain parameter estimates:

$$\psi_k' = \arg_{\psi_k'} \max p(\psi_k' | \mathbf{X}_k, \Lambda)$$

$$\mathbf{a}_k = (\mathbf{G}_k^T \mathbf{G}_k)^{-1} \mathbf{G}_k^T \mathbf{x}_k$$

6. Resample $[\psi_k^{(i)}] = \text{RESAMPLE}[\{\psi_k', w_k^{(i)}\}_{i=1}^N]$

7. Reversible jump move

For $i = 1 : N$

- Sample $u \sim \mathcal{U}_{[0,1]}$; select move type according to probabilities b_k, d_k, u_k .
- Propose new state $\{P_k^{*(i)}, \Theta_k^{*(i)}\}$ and evaluate acceptance probability α .
- Sample $u \sim \mathcal{U}_{[0,1]}$; if $u \leq \alpha$, accept new state $\{P_k^{*(i)}, \Theta_k^{*(i)}\}$, otherwise retain old state $\{P_k^{(i)}, \Theta_k^{(i)}\}$.

which possess significant variety, the added flexibility of the SoLC model is useful. We will then show that the algorithm is able to detect changes in the number of subcomponents when using frequency modulated signals. We will finally present some results from tests on field-recorded bat echolocation calls.

5.1 SoS Model vs SoLC Model

As a test signal, we use a quadratic chirp having a constant-frequency (CF) tail (figure 3), at an SNR of approximately 12dB. Using a normalised sampling frequency of 1 Hz, a spectrogram of the signal (without noise) highlights the presence of multiple peaks (multiple dark tracks), which can lead to over-modelling of the signal. The number of components in the simulation is consequently constrained to 1, and 100 particles are used for all the simulations. The window length is set to 129 samples. 500 Monte Carlo runs were performed for each model and the average MSE was plotted.

A comparison of the average MSE shows that the SoLC model outperforms the SoS model for the FM section of the signal, and is only slightly worse in the CF tail. Also compared with these models, is the MSE from tracking peaks in the spectrogram. The method is similar, but not identical to the method of Dubois *et al* [2], relying implicitly on a SoS model, and is found to perform about as well as the SoS model used here. The use of a hamming window instead of a rectangular window, which is used in all the simulations here, can provide a slight improvement for the STFT tracker since this lowers the sidelobes in the Fourier transform.

The added complexity from the SoLC model allows us to obtain significantly better estimates from chirp signals. The presence of multiple peaks in the Fourier transform will result in over-modelling in the case of the SoS model, however, the SoLC model is flexible enough to deal with this non-stationarity.

5.2 Frequency Tracking Results

The purpose of these results is to illustrate filter performance at points where signals start and end. The algorithm is able to track multiple components, although there is some uncertainty in the estimate at the start and end of signals. At crossing points, the al-

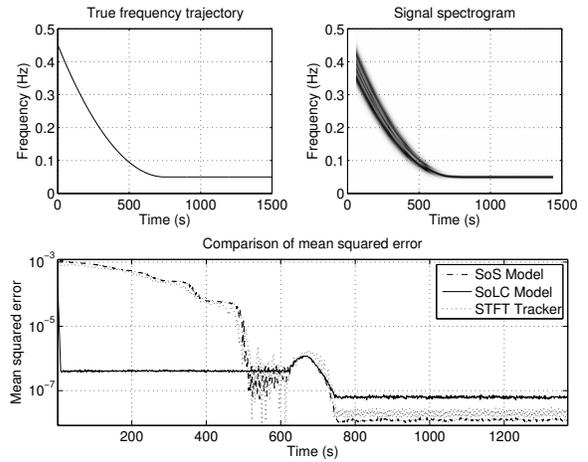


Figure 3: Comparison of SoS and SoLC models

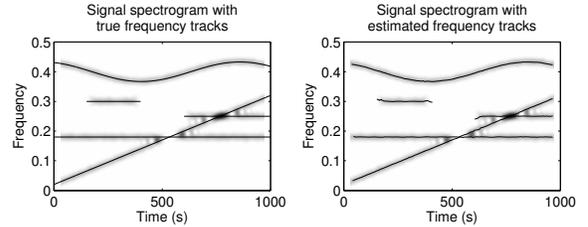


Figure 4: Test using multiple (and crossing) signals

gorithm tracks the separate components and does not approximate them as a single component.

The signal used is similar to that used in [3]. Figure 4 shows the spectrogram of the signal with the true frequency trajectories overlaid. The signal incorporates FM signals, multiple frequency crossings as well as a changing number of components, although none of the components is amplitude-modulated. The SNR of the signal is approximately 20 dB.

In the simulation, the number of particles used is 1000; a window-length of 65 samples is chosen, and the hyperparameters are set as $\alpha_n = \beta_n = 0$, $\alpha_\delta = 1000$, $\beta_\delta = 2$, $\lambda = 1$. The hyperparameters are deliberately chosen so as to specify vague priors on the parameters. Using $(\alpha_\delta = 1000, \beta_\delta = 2)$ specifies a prior distribution on the ‘‘SNR’’ term with mean 30dB.

The presence of discontinuities within the sliding window causes problems for the filter, which tends to over-model the signal in an attempt to fit the discontinuity. To overcome this problem, we assume that two frequency components will be no closer than a predefined limit. This is achieved in practice by inserting nulls into the sampling distribution for the frequency parameter. The disadvantage of this, however, is that frequencies which are very close together will not be detected using the reversible jump sampler. An alternate way of limiting this problem is to use a very short window.

5.3 Frequency Tracking Applied to Bat Calls

Results are shown here from testing the algorithm with two bat echolocation calls recorded in the field. 1000 particles were used and a window-length of 65 samples was specified. A truncated uniform prior is used for the distribution of the frequency component. Since these recordings contain a significant amount of low-frequency noise, the truncated prior is used to disregard the low-frequency noise band. The sampling distribution for adding new frequency components is chosen as a modified truncated prior containing nulls around already existing frequencies.

Figure 5 is a good quality recording of a Pipistrelle bat call (16 bits per sample (bps) with 750 kHz sampling rate). The signal has a single dominant component with another component approximately 30-dB below that. The algorithm is able to track the non-linear chirp as well as a higher harmonic. Figure 6 shows a recording from a

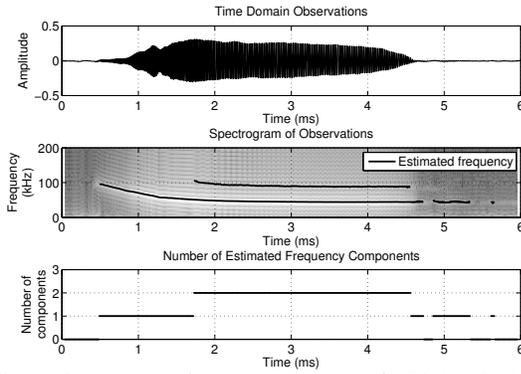


Figure 5: Analysis of echolocation call of a Pipistrelle bat

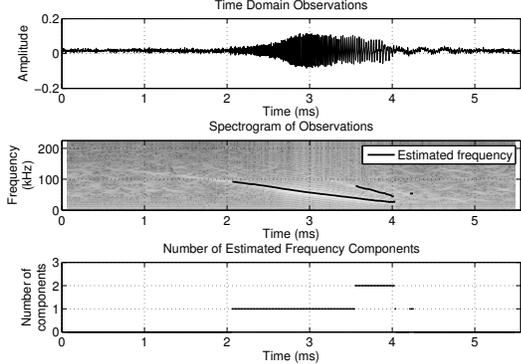


Figure 6: Analysis of echolocation call of a Natterer's bat

Natterer's bat with an SNR of 16 dB (8 bps with 450 kHz sampling rate). The two frequency components are detected in the signal.

The application of the algorithm to bat chirps allows us to extract a set of instantaneous frequencies from the signals. It is possible to transform these frequencies into individual signal tracks (for example, by applying target tracking algorithms). These tracks can then be studied to better describe the nature of bat calls in terms of linear/hyperbolic chirps with relevant parameters.

6. SUMMARY AND CONCLUSIONS

The method described here provides a means for detecting a time-varying number of dynamic frequencies by locally modelling the signal as a sum of linear chirps. The algorithm is based around direct inference from the observations without any form of preprocessing. The use of a sum of linear chirps basis is shown to offer advantages over a sum of sinusoids basis. A particle filtering framework allows frequency tracking in a sequential framework, while the RJMCMC sampler permits detection of the start/end of components. A marginalisation operation is used to reduce the dimension of the parameter state space. This marginalisation of amplitudes and noise variance is shown to be valid for the GLM framework.

A. MARGINALISATION FOR SEQUENTIAL UPDATE

The particle filter weight update equation is written as

$$w_k = w_{k-1} \frac{p(\mathbf{x}_k | \boldsymbol{\psi}_k) p(\boldsymbol{\psi}_k | \boldsymbol{\psi}_{k-1})}{q(\boldsymbol{\psi}_k | \boldsymbol{\Psi}_{k-1}, \mathbf{X}_k)}$$

where w_{k-1} is the weight of a particle at the previous time instant $k-1$. If $\boldsymbol{\psi}'_k$ denotes the state parameters we wish to estimate, and Υ_k is the parameters we wish to marginalise out of the estimation, then the transition prior $p(\boldsymbol{\psi}_k | \boldsymbol{\psi}_{k-1})$ can be rewritten as

$$\begin{aligned} p(\boldsymbol{\psi}_k | \boldsymbol{\psi}_{k-1}) &= p(\boldsymbol{\psi}'_k, \Upsilon_k | \boldsymbol{\psi}'_{k-1}, \Upsilon_{k-1}) \\ &= p(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}, \Upsilon_k, \Upsilon_{k-1}) p(\Upsilon_k | \Upsilon_{k-1}, \boldsymbol{\psi}'_{k-1}) \\ &= p(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}) p(\Upsilon_k | \Upsilon_{k-1}) \end{aligned}$$

We make a further approximation that the Υ parameters are completely independent across blocks, such that $p(\Upsilon_k | \Upsilon_{k-1}) = p(\Upsilon_k)$, where $p(\Upsilon_k)$ represents some prior distribution. This allows us to carry out the marginalisation with relative ease by removing any dependence on previous states.

The sampling distribution reflects the parameters being drawn and is specified as $q(\boldsymbol{\psi}_k | \boldsymbol{\Psi}_{k-1}, \mathbf{X}_k) = q(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}, \mathbf{x}_k)$, i.e., dependent on parameters of interest and the most recent observation.

Evaluation of the weight w'_k is key to computation of the marginalised posterior distribution.

$$w'_k = \int \dots \int \frac{p(\mathbf{x}_k | \boldsymbol{\psi}_k) p(\boldsymbol{\psi}_k | \boldsymbol{\psi}_{k-1}) p(\boldsymbol{\Psi}_{k-1} | \mathbf{X}_{k-1})}{q(\boldsymbol{\psi}_k | \boldsymbol{\Psi}_{k-1}, \mathbf{X}_k) q(\boldsymbol{\Psi}_{k-1} | \mathbf{X}_{k-1})} d\Upsilon_{0:k}$$

The marginalised weight evaluates to:

$$w'_k \propto \frac{p(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1})}{q(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}, \mathbf{x}_k)} \int p(\mathbf{x}_k | \boldsymbol{\psi}_k) p(\Upsilon_k) d\Upsilon_k \quad (14)$$

$$\times \int \dots \int p(\boldsymbol{\Psi}_{k-1} | \mathbf{X}_{k-1}) d\Upsilon_{0:k-1}$$

$$\propto w'_{k-1} \times \frac{p(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1})}{q(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}, \mathbf{x}_k)} \int p(\mathbf{x}_k | \boldsymbol{\psi}_k) p(\Upsilon_k) d\Upsilon_k \quad (15)$$

$$\propto w'_{k-1} \times \frac{p(\mathbf{x}_k | \boldsymbol{\psi}'_k) p(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1})}{q(\boldsymbol{\psi}'_k | \boldsymbol{\psi}'_{k-1}, \mathbf{x}_k)} \quad (16)$$

The integral over $d\Upsilon_k$ in equation 15 is easily performed for the GLM when the likelihood function resembles equation 5 [7].

REFERENCES

- [1] J. A. Simmons, M. B. Fenton, and M. J. O'Farrell, "Echolocation and pursuit of prey by bats," *Science*, vol. 203, Jan 1979.
- [2] C. Dubois, M. Davy, and J. Idier, "Tracking of time-frequency components using particle filtering," *Proceedings (ICASSP '05) IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 4, Mar 2005.
- [3] C. Andrieu, M. Davy, and A. Doucet, "Efficient particle filtering for jump Markov systems. Application to time-varying autoregressions," *IEEE Transactions on Signal Processing*, vol. 51, no. 7, Jul 2003.
- [4] C. Dubois and M. Davy, "Joint detection and tracking of time-varying harmonic components: A flexible Bayesian approach," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 15, no. 4, May 2007.
- [5] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artech House Publishers, 2004.
- [6] G. L. Bretthorst, *Bayesian Spectrum Analysis and Parameter Estimation*, Lecture Notes in Statistics. Springer-Verlag, 1988.
- [7] J. J. K. ÓRuanaidh and W. Fitzgerald, *Numerical Bayesian Methods Applied to Signal Processing*, Springer, 1996.
- [8] P. J. Green, "Reversible jump Markov chain Monte Carlo computation and Bayesian model determination," *Biometrika*, vol. 82, no. 4, 1995.
- [9] C. Andrieu and A. Doucet, "Joint Bayesian model selection and estimation of noisy sinusoids via reversible jump MCMC," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, Oct 1999.
- [10] A. Jasra, A. Doucet, D. A. Stephens, and C. C. Holmes, "Interacting sequential Monte Carlo samplers for trans-dimensional simulation," *Computational Statistics & Data Analysis*, vol. 52, no. 4, 2008.
- [11] J.-R. Larocque, J. P. Reilly, and W. Ng, "Particle filters for tracking an unknown number of sources," *IEEE Transactions on Signal Processing*, vol. 50, no. 12, Dec 2002.