



# Development of Radar Altimetry Data Processing in the Oceanic Coastal Zone

*ESA/ESRIN Contract No. 21201/08/I-LG*

**EWP3 - Deliverable D3.2**

## **Bayes Linear Retracking of Radar Altimetry Data**

**Peter Challenor**

VERSION 1.1 28 May 2012

©The Copyright of this document is the property of National Oceanography Centre (NOC). It is supplied on the express terms that it be treated as confidential, and may not be copied, or disclosed, to any third party, except as defined in the contract, or unless authorised by NOC in writing.

National Oceanography Centre European Way, Southampton SO14 3ZH,  
United Kingdom Tel: +44 (0)23 80596404 Fax: +44 (0)23 80596400

**Code** COASTALT2-D32-11    **Edition** 1.1    **Date** 28 May 2012

**Client** European Space Agency    **Final User**

	Name	Signature	Date
<b>Written by</b>	Peter Challenor		18/5/2012
<b>Checked by</b>	Paolo Cipollini		28/5/2012

<b>DISTRIBUTION</b>	<b>Affiliation</b>
Jérôme Benveniste, Salvatore Dinardo, Bruno Lucas	ESA
Paolo Cipollini, Peter Challenor, Helen Snaith, Luke West	NOC
Stefano Vignudelli	CNR
Jesus Gomez-Enri	Univ. Cadiz
Cristina Martin	STARLAB
Joana Fernandes, Alexandra Nunes	Univ. Porto
Susana Barbosa	Univ. Lisbon

## Revision History

Issue	Date	Change
1	28 May 2011	Initial release
1.1	28 May 2012	Response to ESA comments



## Contents

<b>1</b>	<b>Background</b>	<b>4</b>
1.1	Introduction . . . . .	4
1.2	Altimeter retracking . . . . .	4
1.3	Bayesian Retracking . . . . .	5
<b>2</b>	<b>The Bayes Linear Retracker</b>	<b>6</b>
<b>3</b>	<b>Further Implementation Details</b>	<b>7</b>
<b>4</b>	<b>A Simple Example</b>	<b>8</b>
<b>5</b>	<b>Numerical Issues</b>	<b>10</b>
<b>6</b>	<b>Application in the Coastal Zone</b>	<b>12</b>
<b>7</b>	<b>Comments</b>	<b>12</b>
<b>8</b>	<b>Acronyms</b>	<b>13</b>
<b>A</b>	<b>Installing the software</b>	<b>14</b>
<b>B</b>	<b>Example R code</b>	<b>15</b>

## 1 Background

### 1.1 Introduction

Radar altimeter retracking is essentially a problem of statistical estimation. We have a noisy altimeter waveform and we wish to estimate some geophysical parameters from it. There are two ways of addressing this problem. There are *non-parametric* retrackers that estimate the tracking point on the Earth's surface without a model of the form of the reflection. Such retrackers are robust and are usually used over topographic surfaces such as land or ice. Over the ocean we have parametric models of what we expect the return to look like, for example the Brown model or non-linear variants of it. If we have such models we can use *parametric* statistical methods to estimate the geophysical parameters. This is what is commonly done for altimeter retracking over the oceans. For example, in the Envisat RA-2 a non-parametric retracker (OCOG) is used to track the altimeter signal on board the satellite and a Brown type model is then used in a parametric retracker on the ground. Most current parametric retrackers consider each waveform separately and use either maximum likelihood or least squares to perform the estimation. One exception to the rule of analysing each waveform separately is Sandwell and Smith, [1]. They perform a two pass estimation. In the first pass they carry out a conventional waveform by waveform estimation. They then smooth their significant wave height estimates along the altimeter track. They then retrack assuming that the significant waveheight is known. The estimates of sea surface height obtained by this method have much less uncertainty on them. This is not surprising, as by assuming in the second stage that we *know* significant wave height we omit one of the sources of uncertainty. Thus the reduction of uncertainty arises partly from ignoring some of the input uncertainty and is thus to some extent spurious.

### 1.2 Altimeter retracking

The radar altimeter emits a narrow pulse of radar energy which is reflected from the sea surface and received back at the satellite. This pulse is modified in shape by its interaction with the sea surface. Reflections from the crests of the waves return earlier than the mean sea surface whereas reflections from wave troughs come back later. It should be noted that this description is a simplification of how altimeters work. In fact the pulse is a chirp in the frequency domain and the waveform we analyse is the Fourier transform of the radar signal received by the satellite.

The form of the return can be described analytically as  $W(\tau, \theta)$  where  $\tau$  is the number of the gate in the radar receiver and  $\theta$  is a vector of parameters affecting the form of the waveform. This is the expected form of the return as it is subject to two forms of uncertainty. First there is the *thermal* noise. This is signal that would be received if we transmitted nothing and is the background radiation at the radar frequency. It can be described as white (but non-Normal) noise within the pulse and is independent from pulse to pulse. We denote the thermal noise by  $\epsilon_\tau$ . The second source of noise is the *fading* noise<sup>1</sup>. This is the noise on the actual received signal. For a single pulse this is an exponential distribution with mean and standard deviation both given by  $W(\tau, \theta)$ . To reduce the variance, a number of pulses ( $n$ ) are averaged to give a single mean

<sup>1</sup>Fading noise is also known as speckle noise

waveform  $\bar{W}_n(\tau, \theta)$ . Assuming the fading noise is independent from waveform to waveform this has a  $\Gamma$ -distribution with mean  $W(\tau, \theta)$  and standard deviation  $W(\tau, \theta)/n$ .

The standard way of retracking an altimeter waveform is to use either maximum likelihood or least squares. Both can be expressed as the solution to the equation

$$\operatorname{argmin}_{\theta} \sum_{\tau=0}^{n_{gates}} w_{\tau} (x_{\tau} - W(\tau, \theta))^2 \quad (1)$$

where  $x_{\tau}$  is the value of the waveform at time  $\tau$ ,  $w_{\tau}$  is a weight.  $w_{\tau} = 1$  for least squares and  $= 1/W(\tau, \theta)$  for maximum likelihood., and  $\operatorname{argmin}$  means the arguments, here  $\theta$ , which minimise the function.

The problem with this method of retracking is that it ignores our knowledge about the relationship between adjacent waveforms. We know that adjacent measurements will have similar values of the geophysical parameters yet we ignore this and analyse our data as if the order were irrelevant.

One alternative to the independent tracking of individual waveforms has been produced by (author?) [1]. In this paper the waveforms are initially tracked in a conventional way as described above. A smooth function is then fitted to the significant wave height values and the waveforms are then retracked but this time assuming that the wave height is known. This produces estimates of sea surface height that are much smoother with less noise on them than the initial retracked data. Although this method is attractive as it does take into account at least for wave height the sequential nature of the waveforms it underestimates the uncertainty on the final estimate of sea surface height. This is because in the second retracking of the data the uncertainty on the estimate of significant wave height is ignored. This could be restored by sampling from the smooth functions that are used to fit the significant wave height values but as far as we know this has never been tried.

### 1.3 Bayesian Retracking

Our aim is to include information from other waveforms in the retracking of the  $i^{th}$ . We want to do this without inducing spurious correlation that would come from conventional smoothing where we perform a weighted average along track. One solution to this problem is to use Bayes theorem.

From Bayes theorem we have

$$\pi(\theta|w) \propto \pi(w|\theta)\pi(\theta) \quad (2)$$

Where  $\pi(\theta|w)$  is the posterior pdf of  $\theta$  given the measured waveform  $w$ ,  $\pi(w|\theta)$  is the likelihood of the measured waveform and  $\pi(\theta)$  is the prior pdf of  $\theta$ . As in most applications of Bayes theorem we can calculate the constant of proportionality because the posterior probability density ( $\pi(\theta|w)$ ) integrates to 1. Our prior information ( $\pi(\theta)$ ) is derived mainly from the re-tracking of previous waveforms but could also be derived from other information. For simplicity we will assume here that we retrack waveforms sequentially, this is not necessary for the theory though.

Our procedure for retracking is therefore as follows. We have some initial belief in what the values of the parameters governing the shape of our waveform is. This is the prior pdf,  $\pi(\theta)$ .

Normally this will come from the re-tracking of the previous waveform, but with the variance increased to account for the fact that we expect the values to have changed from that waveform. We now input this  $\pi(\theta)$  and the likelihood of the observed waveform ( $\pi(w|\theta)$ ) into equation 2 and solve for  $\pi(\theta|w)$ . We can then obtain point estimates for  $\theta$  by taking the mode or mean of the posterior pdf. Equation 2 cannot be solved analytically but can be solved using MCMC (Monte Carlo Markov Chain) methods or by sequential Monte Carlo (particle filters) or by a combination [2]. A full Bayesian analysis using MCMC methodology is presented in [3] Although this gives a solution it is computationally extremely expensive and is not suitable for operational use. Most of that computational effort is used to sample the full pdfs of the geophysical parameters. These pdfs are of limited use and would probably be ignored by the user community. Currently altimeter products do not even report the variance matrix of the estimates even though this is easily calculated as part of the fitting process.

## 2 The Bayes Linear Retracker

The full Bayesian solution presented in the last section calculated the full probability density function of the parameters of the waveform and we pointed out that the full pdf is not required. There is an alternative form of Bayesian statistics, Bayes linear, that instead of calculating the full pdf works with only the first two moments of the distributions: the means, variances and covariances. For details of Bayes linear statistics see [4]. Instead of updating our beliefs using Bayes theorem (equation 1) we use a Bayes linear update. In our case this is given by

$$E(\theta|w) = E(\theta) + Cov(\theta, w)V(w)^{-1}(w - E(w(\theta))) \quad (3)$$

for the adjusted expectation and by

$$V(\theta|w) = V(\theta) - Cov(\theta, w)V(w)^{-1}Cov(w, \theta) \quad (4)$$

for the adjusted variance matrix.

Let us consider each of these terms one by one.

$E(\theta)$  is our prior expectation for the geophysical parameters.

$w$  is the measured waveform as a vector. We denote the value at bin  $\tau$  with  $w(\tau)$ .

$V(w)$  is the variance matrix of the waveform. Normally we assume that the individual bins are uncorrelated so this is a diagonal matrix with diagonal terms given by either  $w/n$  or  $E(w)/n$  (where  $n$  is the number of pulses averaged in creating the waveform we are tracking).

$Cov(\theta, w)$  is the covariance matrix between the  $\theta$  and the waveform values.

$E(w(\theta))$  is the expected value of the waveform  $w$ .

Let us consider the interpretation of each term in (3). The first term is easy to interpret. It is our prior expectation of the values of the parameters or in other words our belief about the value of the parameters before we collect the waveform. The second term is how these beliefs are modified by the data. This term has three parts. Starting from the right we have  $(w - E(w(\theta)))$ . This is the difference between the measured waveform ( $w$ ) and our expectation of the waveform ( $E(w(\theta))$ ). This difference is weighted by its variance  $V(w(\theta))$ . This weighted difference is in the space of the waveform so we need to convert this into a correction in  $\theta$  space. This is done

by multiplying by the  $Cov(\theta, \mathbf{w})$ , which can be thought as the amount that  $\theta$  changes for a given change in  $\mathbf{w}$ .

Two terms,  $E(\mathbf{w})$  and the  $Cov(\theta, \mathbf{w})$ , need to be further expanded. We do this using a scaled Taylor series expansion. For details see the appendix.

The advantages of the Bayes linear retracker compared to MLE or least squares are:

- Fast
  - There is no non-linear optimisation involved simply matrix algebra.
- General
  - The method works with any parametric form for a waveform including those used in the coastal areas, where specular reflections from ‘bright’ targets may be included.
- Gives both estimates and uncertainty on the geophysical parameters
- Easy to construct statistical tests to distinguish between different waveform models
  - This is particularly important in coastal analysis where we will want to test for the presence of ‘bright’ targets.
- Sequential
  - Bayesian form means that it is naturally sequential

The next two sections of this report give some examples of the use of a Bayes Linear retracker and describe an R package that implements it.

### 3 Further Implementation Details

Because  $\mathbf{w}$  is a non-linear function  $E(\mathbf{w}(\theta)) \neq \mathbf{w}(E(\theta))$ . To calculate  $E(\mathbf{w}(\theta))$  we expand  $\mathbf{w}(\theta)$  as a Taylor series around  $E(\theta)$ .

$$\mathbf{w}(\theta) = \mathbf{w}(E(\theta)) + \sum_{i=1}^n \frac{\partial \mathbf{w}(E(\theta))}{\partial \theta_i} (\theta_i - E(\theta_i)) + \sum_{i=1}^n \sum_{j=1}^n \frac{\partial^2 \mathbf{w}(E(\theta))}{\partial \theta_i \partial \theta_j} \frac{(\theta_i - E(\theta_i))(\theta_j - E(\theta_j))}{2} + \dots \quad (5)$$

If we denote the gradient of a vector by  $D$  and the Hessian by  $D^2$  we can express this as a matrix equation

$$\mathbf{w}(\theta) = \mathbf{w}(E(\theta)) + (\theta - E(\theta))^T D(\mathbf{w}(\theta), \theta) + \frac{(\theta - E(\theta))^T D^2(\mathbf{w}(\theta), \theta) (\theta - E(\theta))}{2} + \dots \quad (6)$$

We are only interesting in first and second moments so we drop the higher order terms. Now take expectations

$$E(\mathbf{w}(\theta)) = E[\mathbf{w}(E(\theta)) + (\theta - E(\theta))^T D(\mathbf{w}(\theta), \theta) + \frac{(\theta - E(\theta))^T D^2(\mathbf{w}(\theta), \theta) (\theta - E(\theta))}{2}] \quad (7)$$

Since  $E(\theta - E(\theta)) = 0$  and  $E[(\theta - E(\theta))^T(\theta - E(\theta))] = V(\theta)$  we have

$$E(w(\theta)) = w(E(\theta)) + \left( \frac{D^2(\mathbf{w}(\theta), \theta)V(\theta)}{2} \right) \quad (8)$$

Now consider the  $Cov(\theta, \mathbf{w}(\theta))$ . Again we use the Taylor series expansion for  $\mathbf{w}$ .

$$Cov(\theta, \mathbf{w}(\theta)) = E(\theta \mathbf{w}(\theta)^T) - E(\theta)E(\mathbf{w}(\theta))^T \quad (9)$$

$$E(\theta \mathbf{w}(\theta)^T) = E[\theta \mathbf{w}(E(\theta))^T + \theta (D(\mathbf{w}(\theta), \theta))^T (\theta - E(\theta))^T] \quad (10)$$

$$\begin{aligned} & + \theta \left( \frac{(\theta - E(\theta))^T D^2(\mathbf{w}(\theta), \theta) (\theta - E(\theta))}{2} \right)^T \\ & = E(\theta) \mathbf{w}(E(\theta))^T + E(\theta(\theta - E(\theta))D(\mathbf{w}(\theta), \theta))^T + \text{skewness terms} \end{aligned} \quad (11)$$

Dropping the skewness terms and noting that

$$E(\theta(\theta - E(\theta))) = V(\theta) \quad (12)$$

we have

$$Cov(\theta, \mathbf{w}(\theta)) = E(\theta) \mathbf{w}(E(\theta))^T + V(\theta)D(\mathbf{w})^T - E(\theta) (E(\mathbf{w}(\theta)^T)) \quad (13)$$

Replacing  $E(\mathbf{w}(\theta))$  with (8) we get

$$\begin{aligned} Cov(\theta, \mathbf{w}(\theta)) &= E(\theta) \mathbf{w}(E(\theta))^T + V(\theta)D(\mathbf{w})^T \\ &- \left( E(\theta) \mathbf{w}(E(\theta))^T - E(\theta) \left( \frac{D^2(\mathbf{w}(\theta), \theta)V(\theta)}{2} \right)^T \right) \end{aligned} \quad (14)$$

$$Cov(\theta, \mathbf{w}(\theta)) = V(\theta)D(\mathbf{w})^T - E(\theta) \left( \frac{D^2(\mathbf{w}(\theta), \theta)V(\theta)}{2} \right)^T \quad (15)$$

Thus we can approximate both  $E(w(\theta))$  and  $Cov(\theta, w(\theta))$ .

## 4 A Simple Example

For an illustrative example we will use a simplified form of the full Brown model. This has four parameters:  $\sigma_0$  (the amplitude of the waveform),  $h$  (the position of the waveform in the window  $\equiv$  the sea surface height),  $H_s$  (the slope of the leading edge  $\equiv$  the significant wave height) and  $t_n$ , the thermal noise. Our simplified model is given by

$$E(w(\tau)) = \frac{\sigma_0}{2} \left( 1 + erf \left( \frac{\tau - h}{h_s} \right) \right) + t_n \quad (16)$$

$erf()$  is defined as in [5]

$$erf(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$



The  $V(w(\tau))$  is equal to  $E(w(\tau))/n_p$  where  $n_p$  is the number of radar pulses averaged into a single waveform. Giving  $w(\tau)$  a  $\chi^2$  distribution with  $n_p$  degrees of freedom we can generate waveforms as shown in figure 1.

For the Taylor series approximation to the moments (see the appendix) we need to have the gradient and Hessian of the waveform w.r.t. each of the parameters. For a particular value of  $\tau$  these are given by

$$\frac{\partial w(\tau)}{\partial t_n} = 1 \quad (17)$$

$$\frac{\partial w(\tau)}{\partial \sigma_0} = \frac{1}{2} \left( 1 + \operatorname{erf} \left( \frac{\tau - h}{h_s} \right) \right) \quad (18)$$

$$\frac{\partial w(\tau)}{\partial h} = \frac{\sigma_0 e^{-\frac{(\tau-h)^2}{h_s^2}}}{\sqrt{\pi} h_s} \quad (19)$$

$$\frac{\partial w(\tau)}{\partial h_s} = \frac{\sigma_0 (\tau - h) e^{-\frac{(\tau-h)^2}{h_s^2}}}{\sqrt{\pi} h_s^2} \quad (20)$$

$$\frac{\partial^2 w(\tau)}{\partial t_n^2} = 0 \quad (21)$$

$$\frac{\partial^2 w(\tau)}{\partial \sigma_0^2} = 0 \quad (22)$$

$$\frac{\partial^2 w(\tau)}{\partial h^2} = \frac{2\sigma_0 (\tau - h) e^{-\frac{(\tau-h)^2}{h_s^2}}}{\sqrt{\pi} h_s^3} \quad (23)$$

$$\frac{\partial^2 w(\tau)}{\partial h_s^2} = \frac{2\sigma_0 (\tau - h) e^{-\frac{(\tau-h)^2}{h_s^2}}}{\sqrt{\pi} h_s^3} \left[ 1 - \frac{(\tau - h)^2}{h_s^2} \right] \quad (24)$$

$$\frac{\partial^2 w(\tau)}{\partial h_s \partial t_n} = \frac{\partial^2 w}{\partial h \partial t_n} = \frac{\partial^2 w}{\partial \sigma_0 \partial t_n} = 0 \quad (25)$$

$$\frac{\partial^2 w(\tau)}{\partial h \partial \sigma_0} = \frac{e^{-\frac{(\tau-h)^2}{h_s^2}}}{h_s \sqrt{\pi}} \quad (26)$$

$$\frac{\partial^2 w(\tau)}{\partial h_s \partial \sigma_0} = -\frac{e^{-\frac{(\tau-h)^2}{h_s^2}} (\tau - h)}{h_s^2 \sqrt{\pi}} \quad (27)$$

$$\frac{\partial^2 w(\tau)}{\partial h_s \partial h} = \frac{\sigma_0 e^{-\frac{(\tau-h)^2}{h_s^2}}}{h_s^2 \sqrt{\pi}} \left( 1 - \frac{2(\tau - h)^2}{h_s^2} \right) \quad (28)$$

As an example we simulate a waveform with the following parameters:  $h=1.1$ ,  $h_s=4.1$ ,  $\sigma_0=10.5$ ,  $t_n=1.0e-06$  with 64 gates each length 1 and averaged over 100 independent individual echoes. For numerical stability we need all our parameters to have approximately the same order so we rescale the thermal noise ( $t_n$ ) to be one.

Parameter	$h$	$h_s$	$\sigma^0$	$t_n$
Truth	1.10	4.10	10.50	1.00e-06
MLE	1.10	4.10	10.49	0.0096
Bayes Linear	1.17	4.05	10.29	0.96

Table 1: The estimates of the geophysical parameters from the MLE and Bayes linear (BL) retracker compared to the true values

Parameter	$h$	$h_s$	$\sigma^0$	$t_n$
$h$	0.019	0.010	-0.017	2.28e-06
$h_s$	0.010	0.007	-0.005	2.30e-06
$\sigma^0$	-1.17	-0.005	0.011	-9.03e-07
$t_n$	2.28e-06	2.30e-06	-9.03e-07	4.22e-08

Table 2: Estimates of the variance-covariance matrix for the MLE retracker

The simulated waveform is shown in figure 1. We now have to generate a prior mean and variance matrix. As the mean we use the MLE estimates for the waveform and for the prior variance matrix we use the inverse of the Fisher Information matrix. These are the asymptotic variances for the MLE's [6].

To carry out this analysis we use the following R code set out in appendix B.

Table 1 shows the results of MLE and Bayes linear estimation.

Table 2 shows the MLE variance matrix and Table 3 the Bayes Linear posterior variance matrix. The variances are all slightly smaller than the MLE because of the inclusion of prior information about the parameters and their variances. As tracking proceeds through a series of waveforms this effect will increase as the retracker 'learns' about the sea surface.

## 5 Numerical Issues

Although the method described above has been successfully used to retrace single pulses we have not yet produced code that can robustly retrack a series of pulses and 'learn' about the waveforms in a manner that would clearly demonstrate the advantages of the Bayes Linear methodology. Indeed even the single pulse code will occasionally fail, producing a negative variance. The reason for these problems appears to be the modelling of the thermal noise. The size of the thermal noise, and in particular the magnitude of its variance, means that the eigenval-

Parameter	$h$	$h_s$	$\sigma^0$	$t_n$
$h$	0.015	0.012	-0.017	-4.97e-07
$h_s$	0.012	0.006	-0.005	3.16e-06
$\sigma^0$	-0.017	-0.005	0.011	-4.95e-06
$t_n$	-4.97e-07	3.16e-06	-4.95e-06	4.03e-08

Table 3: Estimates of the variance-covariance matrix for the Bayes Linear retracker




Figure 1: The sample waveform used in the example.

ues of the matrix vary by at least 10 orders of magnitude. Even doing the calculation in double precision does not allow floating point calculations to retain this level of precision. Our attempts to solve this problem by rescaling the variables have failed. The problem is more basic than that. This problem is not confined to the Bayes Linear methodology. Many implementations of MLE do not estimate the thermal noise as part of the MLE calculations. A possible solutions to this problem is what in essence is done in such MLE implementations. A separate estimate is made of the thermal noise, by identifying the part of the waveform before the returned echo and this is then ‘plugged in’ to the Bayes Linear re-tracker. Care of course needs to be taken to avoid the side lobe from the Fourier processing at the start of the echo. Unfortunately we have not yet had time to implement this but it is our intention to do so in future.

It is also possible that the numerical problems derive in part from our Taylor series approximation to derive the correlations between the parameters and the waveform. This is discussed further in the next section.

## 6 Application in the Coastal Zone

Our simple example has been based on a simplified version of the Brown model, an open ocean retracker. The theory however is applicable to any waveform which can be described with an analytical or numerical expression governed by a small number of parameters. For example one model for returns in the coastal zone is an open ocean Brown return plus the sum of a number of specular reflections. If the specular reflections are modelled as narrow Gaussian peaks (from experience this is not an unreasonable assumption) then our waveform can be described thus

$$W(\tau, \theta) = B(\tau, h, H_s, \sigma^0) + \sum_{i=1}^n \alpha_i \exp(-\gamma_i(\tau - \eta_i)^2) \quad (29)$$

where  $B(\cdot)$  is the Brown model,  $\alpha_i$ ,  $\eta_i$  and  $\gamma_i$  are the amplitude, location and width of the  $i^{th}$  specular reflection. Bayes Linear methods cannot only be used to retrack such a waveform but also to test which value of  $n$  gives the ‘best’ fit, for details see [4]. This is another advantage for the Bayes Linear retracker, it is possible to build a likelihood ratio test in conjunction with an MLE retracker but the mathematics is much more complex.

## 7 Comments

In our simple example we used a Taylor series expansion to produce the covariances needed. This required analytical derivatives. For our simple example deriving these derivatives is fairly easy, for more complex waveform models this may not be so simple. In such cases numerical derivatives should work, however with computer algebra packages such as Maple and Mathematica now commonly available obtaining analytical derivatives is now longer the chore it used to be. Although the derivatives for a simple model were derived by hand they were checked using Mathematica.

The Taylor series method is not integral to the Bayes Linear retracker but is simply one way of implementing it. It would be completely impracticable to use Taylor series in some cir-

cumstances. For example with advanced altimeter designs such as delay doppler altimeters the expected waveform may only be available numerically rather than analytically. In such a case it would be impossible to use a Taylor series expansion. An alternative is to fit an emulator to the waveform and use the emulator to calculate the relevant covariance matrices. For details on Bayes Linear emulators see the MUCM website [www.mucm.ac.uk](http://www.mucm.ac.uk). Even for analytical waveforms emulators may be more efficient than Taylor series although we have not investigated this approach yet.

## 8 Acronyms

CRAN - Comprehensive R Archive Network  
MCMC - Markov Chain Monte Carlo  
MLE - Maximum Likelihood Estimation  
MUCM - Managing Uncertainty in Complex Models  
OCOg - Offset Centre of Gravity  
pdf - Probability Density Function

## References

- [1] D.T. Sandwell and W.H.F. Smith. Retracking ERS-1 altimeter waveforms for optimal gravity field recovery. *Geophysical Journal International*, 163(1):79–89, Jan 2005.
- [2] C. Andrieu, A. Doucet, and R. Holenstein. Particle Markov Chain Monte Carlo methods. *J R Stat Soc B*, 72:269–342, 2010.
- [3] J. Severini, C. Mailhes, P. Thibaut, and J.-Y. Tournet. Bayesian estimation of altimeter echo parameters. *IEEE International Geoscience and Remote Sensing Symposium (IGARSS) 2008*, 3:238–241, 2008.
- [4] M. Goldstein and D. Wooff. *Bayes Linear Statistics: Theory and Methods*. Wiley, Chichester, 2007.
- [5] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions*. National Bureau of Standards, 1964.
- [6] P.G. Challenor and M.A. Srokosz. The extraction of geophysical parameters from radar altimeter return from a nonlinear sea surface. In S.R. Brooks, editor, *Mathematics in Remote Sensing*, pages 257–268. Oxford University Press, 1989.

## A Installing the software

The R package *alttracker* with the implementation of the Bayes Linear Retracker described in the present document, as well as a manual describing its usage are available from the COASTALT FTP site under:

```
/utilities/bayes_linear_retracker/alttracker_R
```

First R must be installed. R is free and can be downloaded from the CRAN (Comprehensive R Archive Network) (<http://cran.r-project.org/>) or one of the many mirrors around the world. There are executables available for Windows and MacOSX and some flavours of UNIX. Other systems need to compile it but comprehensive instructions are available, and a very helpful r-help e-mail list. The *alttracker* package does not require any additional R packages to be installed.

The *alttracker* package is installed in the same way as any other R package. First the *alttracker* package file must be downloaded. This is a .tgz (or .tar.gz) file and should not be decompressed. In some R installations (e.g. MacOSX) this can be done via a menu but the general command on UNIX is

```
R CMD INSTALL -l /path/to/library alttracker
```

where */path/to/library* is the path to the directory containing *alttracker.tgz*. This should also work on Windows but reference should be made to the R manual 'R Installation and Administration' available from CRAN.

Once *alttracker* has been installed it still has to be loaded into the R session before it can be used. This is done with the R command *library(alttracker)*. See the example session in Appendix B.

## B Example R code

```
library(alttracker)
# set up the array of gate timings
tarray =setuptau(64,3)
# the true parameter values
h=1.1
hs=4.1
sigma0=10.5
tn=1.0e-6
# and scaling values
scaling = c(1,1,1,1)
# Now simulate a waveform
wavesim = wavesamp(h,hs,sigma0,tn,tarray,nsamp=100,scaling=scaling)
# plot the waveform
plot(tarray,wavesim)
# start values for MLE estimation
start =c(h,hs,sigma0,tn)
# setup prior using MLE tracking
prior = setup.prior(as.matrix(wavesim),tarray,start,nsamp=100,scaling)
# rescale to avoid numerical problems
scaling = c(1,1,1,prior$mean[4])
# rescale the prior
invscale = 1/scaling
prior$mean = prior$mean*invscale
prior$var = kronecker(invscale,t(invscale))*prior$var
# the Bayes Linear retracker
post=bayes.linear.retracker(prior$mean,wavesim,prior$var,tarray,100,scaling)
# undo the scaling
postt$mean = post$mean*scaling
post$var = kronecker(scaling,t(scaling))* post$var
```