

Chapter 10

Application to the MH370 Accident

The previous chapters have constructed a Bayesian method for estimating commercial aircraft trajectories using models of the information contained in satellite communications messages and of the aircraft dynamics. This chapter applies the estimator to the accident flight. We show the pdf at the final message time 00:19 and perform several tests that provide confidence in the robustness to the model parameters and measurement characterisation. Probable end of flight dynamics are used to convert the pdf in the air at 00:19 into an ocean search zone. As a further measure of robustness, we show that the model can be initialised at the end of the initial ascent, as was done in the validation flights, and the output pdf remains in the same broad geographic region, even without using the subsequent primary and secondary radar data.

Table 10.1 summarises the available measurement data and the BTO measurement standard deviation, when applicable, since this depends on the message type. The two anomalous R1200 messages and the R600 message occurred during transient phases of operation for the SATCOM equipment so the BFO values reported for these times cannot be used. The BTO measurement variance is also inflated to account for the uncertainty in the correction term.

10.1 The Filter Applied to the Accident Flight

This section summarises the results of applying the variable rate particle filter to the data available for the accident flight. We consider two different treatments of the measurements. Figure 10.1 shows the Google Earth pdf representation for the case where only BTO measurements were used. As with the validation flights, the colour coded paths show the most likely routes to each latitude. Clearly this distribution is multimodal. Based on BTO data, it can be deduced that the aircraft must have manoeuvred in some way between 18:28 and 19:41. However the BTO measurements

Table 10.1 SATCOM messages used for MH370 analysis

Time (UTC)	Measurement type	BTO	BTO std. dev. (μ s)	BFO
18:25:34	Anomalous R1200	Y	43	N*
18:28:05, 18:28:14	R1200	Y	29	Y
18:39:55	C-channel	N	–	Y
19:41:02, 20:41:04, 21:41:26, 22:41:21	R1200	Y	29	Y
23:15:02	C-channel	N	–	Y
00:10:59	R1200	Y	29	Y
00:19:29	R600	Y	63	N*
00:19:37	Anomalous R1200	Y	43	N*

Measurements marked with an asterisk are available but cannot be used as discussed in the text




Fig. 10.1 Set of paths from 18:02 to 00:19 using only BTO measurement weighting (i.e., not using any BFO measurements)

are consistent with two options: heading broadly North or broadly South. Many of the paths turn very early, almost immediately after initialisation. This is because the BTO arc at 18:25 is not consistent with the radial speed prescribed by the initial speed and direction. The 18:22 radar point, while potentially inaccurate, implies that no turn had occurred at that stage and that the speed may have changed. Due to the questions surrounding its accuracy, the filter does not use this radar point, so it postulates early direction changes as well as early speed changes to match the 18:25 BTO. A feature of most of the paths is that they do not choose to make subsequent turns even though there is no BFO data to influence angle likelihood. The BTO data itself from 19:41 onwards is consistent with straight and level flight.




Fig. 10.2 Set of paths from 18:02 to 00:19 using BTO and BFO measurement weighting. *Purple box* shows the November 2015 indicative search region

Figure 10.2 shows how this path set is pruned by including BFO data. The crucial BFO measurements are at 18:28 and 18:39. The first of these is consistent with the same heading as at initialisation, so it implies that the aircraft has not yet turned. In contrast, as discussed in Sect. 5.4, the 18:39 BFO measurement indicates that the aircraft is moving roughly South at this time. Between the two, these measurements restrict the time of the turn to a window between 18:28 and 18:39. This is reflected in the set of surviving paths. Paths that went very far South had to turn earlier and have been rejected as have all of the paths heading to the North towards Asia. The plot also shows the indicative search region as at November 2015, displayed as a purple box.

The main effect of including BFO data is to resolve the ambiguity about the manoeuvre after 18:28. The BFO data does not significantly change the shape of the Southern mode of the pdf. To show this more clearly Fig. 10.3 plots the two pdfs, without and with BFO data, as one dimensional curves parameterised by latitude. The latitude is a nonlinear function of the position around the BTO arc so these pdfs are distorted slightly, but they clearly show the effect of the BFO data: it selects the mode from a multimodal pdf but it does not significantly change the mode shape. Because of this, including only the C-channel BFO measurement has the same effect and changing the assumed BFO noise standard deviation has no significant effect either unless it could be artificially reduced to less than 1 Hz. We conducted several experiments that varied the assumed BFO noise and which BFO values were included. The results are consistent with the curves in Fig. 10.3.




Fig. 10.3 Comparison of the filter pdf without (*top*) and with (*bottom*) BFO measurements

10.2 Manoeuvre Statistics

An interesting question is whether the posterior pdf output from the filter provides any strong estimate of the number and type of manoeuvres after the loss of primary radar coverage. To address this we compare the probability mass of the number of turns for the case of no measurements, i.e., the prior probability from predictions used to create the enormous disc pdf in Fig. 7.4, and the posterior probability after applying the measurements. These are shown in Figs. 7.3 and 10.4 respectively. The posterior distribution does not count turns made between 00:11 and 00:19 because the BFO data at 00:19 is not used (as it is believed to be unreliable); this permits the filter to make superfluous manoeuvres. The histograms were created by counting the number of manoeuvres performed by each particle and then scaling their contribution to the histogram by the particle weight. In the prior probability there is a significant probability of making a very large number of manoeuvres. The particles that contribute to this part of the histogram are those that have selected very short mean manoeuvre times τ .

The histograms show that the posterior number of turns has reduced dramatically, as has the number of speed changes. The likelihood of selecting a sequence of random





Fig. 10.4 Histogram of posterior number of turns and speed changes

speeds that average out to a suitable overall speed is relatively high. Since there is usually at least an hour between BTO arcs, the filter has time to make several speed changes. Only the average of these speeds influences the position and hence the BTO. As discussed several times, the BFO is sufficiently uninformative that the instantaneous velocity at the measurement time is only loosely constrained. However, it is unlikely that a sequence of random turns will “cancel each other out,” and the filter finds very few paths that string together multiple turns when a straight path would suffice. The low support for turns leads to selection of longer mean manoeuvre times, which in turn suppresses superfluous speed changes.

The turn histogram indicates that around half of the paths made more than a single turn. This would appear to be of interest but is in fact misleading. Figure 10.5 shows the posterior distribution of angle as a function of time. The greyscale image shows the probability density at each time slice; the darkest points are most likely. The solid line shows the mean and one-sigma error bars indicate the standard deviation. Note that the horizontal scale is measurement index, not time, but time values are used to label the measurements to give perspective. The angle plot shows that there are a number of particles that make a very small turn before 18:28 and then follow this with a second turn between 18:28 and 18:39. The result is very similar to making a single turn through the same total angle. Similarly, there are a number of particles that make a turn before 18:39 and then a second very small turn immediately afterwards. Again, this is approximately equivalent to a single turn through the same total angle.

Fig. 10.5 Posterior control angle distribution



The apparent large number of double turns is really just a single turn being broken into two segments. There appears to be very few genuine turns later in the flight.

Figure 10.6 shows the marginal distribution of the Mach number between 18:02 and 00:19. The prior is uniform in Mach between 0.73 and 0.84 but the posterior strongly prefers the higher speed part of this range. This is approximately consistent with the most likely speeds for the aircraft over long durations.




Fig. 10.6 Posterior speed distribution

10.3 Residual Measurement Errors

Another measure of the quality of an estimate is the statistical behaviour of the measurement residual errors. If the estimates were exactly the truth then the residual measurement errors would be the actual measurement noise. So these residuals should be uncorrelated, zero mean, and spread consistently with the measurement noise variance. Figure 10.7 shows the marginal distributions of the residual measurement errors. The filter predicts the BTO and BFO at the measurement times based on the estimated position and velocity and the residual is the difference between these




Fig. 10.7 Residual measurement error

predictions and the measured values. Each sample is weighted according to its overall likelihood and the combination gives the marginals in Fig. 10.7. The horizontal axis is again the measurement index and the vertical axis is residual error, dark points have high probability, and the error bars show the one-sigma limits. The figure also overlays the one-sigma lines for the assumed measurement noise. For BTO this varies with the message type, for BFO it is a fixed ± 7 Hz. There is no BTO measurement for the C channel communications that occurred at 18:39 and 23:14. The BFO has been discarded at 18:25 and 00:19 since these readings are thought to be unreliable. The plots only show residuals when a valid measurement is available.

For the BTO residual there are two larger residuals on the final two measurements. These measurements are actually very close together in time: the first is an R600 message at 00:19:29 and the second is an R1200 message at 00:19:37. The reason why both of these measurements show large residuals is that they are not consistent with each other: the residuals have opposite signs, reflecting that one measurement is longer in range than the prediction and the other is shorter. In the absence of a reason to prefer one over the other, we use both measurements and let the filter find paths that are the best statistical fit. The BFO residuals are statistically consistent with the empirical error model.

10.4 Posterior Distribution of Manoeuvre Time Constant


The inference procedure that has been developed samples the manoeuvre time constant, which is a static parameter. In the case of conventional particle resampling, this is not advisable, as subsequent resampling steps generally leave few distinct values of the static parameter. The recursive resampling process described in Chap. 8 is less prone to this difficulty, as many independent sample paths are retained.

The sample support of the posterior distribution of the manoeuvre time constant is demonstrated in Fig. 10.8. The diagram clearly shows that the distribution is non-degenerate. The figure compares the posterior distribution from the filter, which is shown as bars, with the Jeffreys prior, which is shown as a red line. The Jeffreys prior prefers lower mean manoeuvre times but the data does not support these. The posterior is dominated by longer mean manoeuvre times, with 97 % of the distribution having a mean time between manoeuvres of more than one hour, and 83 % having a mean time more than two hours.

10.5 End of Flight

The output of the particle filter is an estimate of the pdf of the aircraft state at 00:19. The aircraft was still in the air at this time and a model is required to describe the distribution of how it may have descended. This has been primarily the responsibility of the ATSB and the other members of the search team. A discussion of the different

Fig. 10.8 Marginal posterior distribution of manoeuvre time constant, τ . *Red line* shows Jeffrey’s prior, as described in Sect. 7.2.1. *Bars* show histogram of samples weighted by measurement likelihood, on bins of width 0.12h. *Gray line* shows histogram on finer bins of width 0.012h



methods used to model the potential motion is presented in [5]. The model for aircraft motion after 00:19 leads to a prioritisation of the search along and around the final BTO arc.


Quantitatively, the model of descent defines a transition distribution $p(\mathbf{x}_{\text{final}}|\mathbf{x}_K)$, which describes the probability density of the final state given the state at time 00:19. In effect this acts as a kernel to spread the distribution at 00:19 via the expression

$$p(\mathbf{x}_{\text{final}}|\mathbf{Z}_K) = \int p(\mathbf{x}_{\text{final}}|\mathbf{x}_K) \sum_{p=1}^P w_K^p \delta(\mathbf{x}_K - \mathbf{x}_K^p) d\mathbf{x}_K = \sum_{p=1}^P w_K^p p(\mathbf{x}_{\text{final}}|\mathbf{x}_K^p) \tag{10.1}$$

The analysis in [5] leads to a probable scenario where the aircraft ran out of fuel at some time between 00:11 and 00:19. The final satellite communications message could be due to the modem rebooting under auxiliary power. Under this hypothesis, the aircraft engines were already unpowered at 00:19. The spread of the kernel function is then determined by the distance over which the aircraft could have moved, which depends on whether or not the aircraft was under human control during this period.

Flight simulator studies of uncontrolled descents have shown a high likelihood of the aircraft reaching zero altitude within 15 nm of the beginning of descent [5]. However, the beginning of descent is not known. It is possible for the aircraft to have travelled farther, especially if a human was controlling the aircraft. As an indicative kernel, and following advice from the ATSB, a uniform disc of radius 15 nm with a Gaussian drop off with standard deviation 30 nm beyond this was chosen; this represents the accident investigators’ assessment of the likely scenarios. Figure 10.9 shows a radial slice through this kernel function. The kernel was convolved with the particle locations at 00:19 to generate the heat map shown in Fig. 10.10. The black region overlaid on the heat-map shows an indicative bounding box on the geographic

Fig. 10.9 Descent kernel



search region. The figure also shows the 99, 95 and 90 % confidence regions, i.e., the smallest geographic area containing each percentage of the pdf.

10.6 Earlier Initialisation

As a further level of validation on robustness, a separate run of the filter was performed initialising using the final ACARS report. This run is not useful for defining a search region because it ignores important data from the primary and secondary radar. However, the result does illustrate the ability of the method to address flight paths with several large turns. Figure 10.11 shows the pdf derived from the 17:01 initialisation. The pdf is overlaid with the one that is presented earlier in this chapter that uses all of the available information. Two versions of the 17:01 initialisation are presented. The first, which ignores the primary radar data, is shown in red in the upper diagram. The second, shown in the lower diagram, re-weights the paths by treating the 18:02 radar point as a measurement with 1° standard deviation in latitude and longitude. This weighting is rather loose but the samples are limited and none of them are sufficiently close to the radar point to permit a tighter match. The filter has given a similar output even though the initialisation point was heading North and there are two very significant manoeuvres reported by the radar in the first hour. This provides further confidence in the ability of the filter to characterise aircraft motion even in complicated situations. The sequence of turns performed by MH370 between 17:01 and 18:02 is much more dynamic than any experienced in the validation flights. When the radar data was ignored the filter found a cluster of paths that travel a significant distance to the North before turning prior to 18:39. These paths conclude at the Northern edge of the search region because they travel too far to the North early on, in a manner which is inconsistent with the primary radar data. In contrast, the

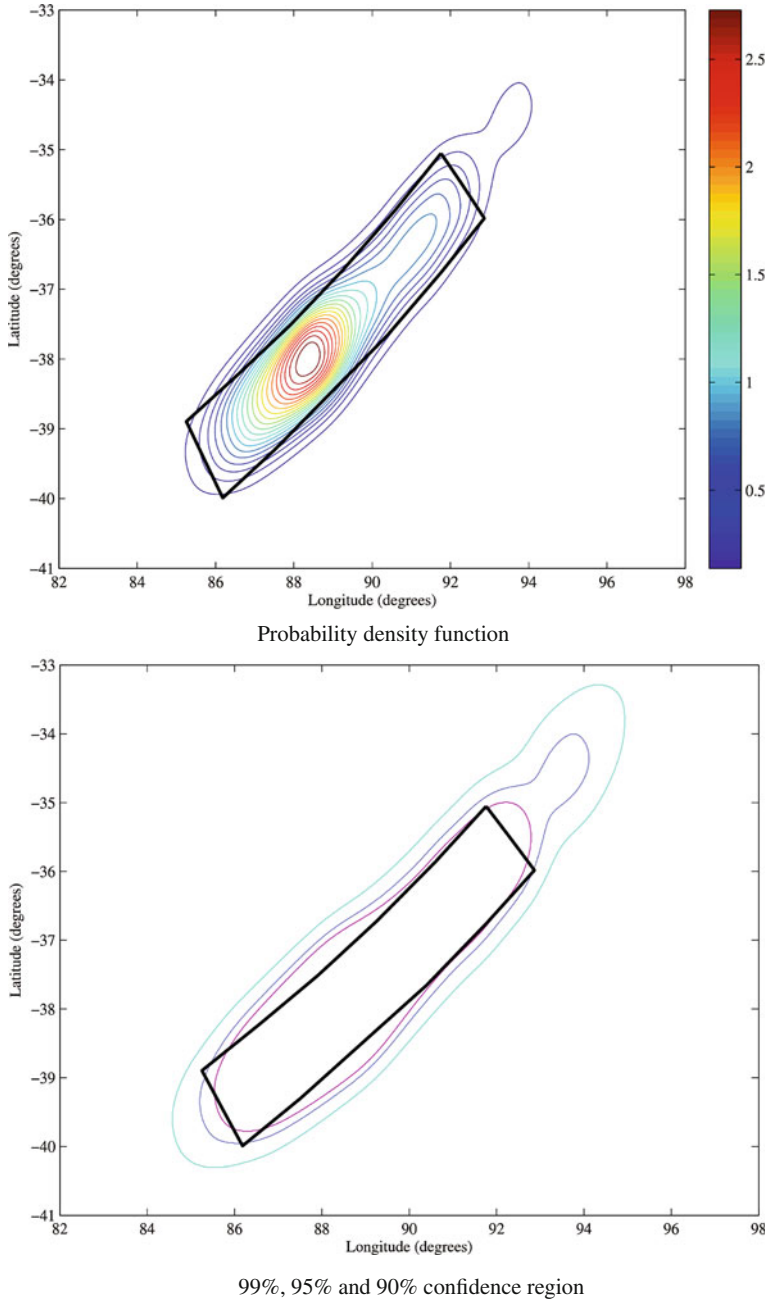




Fig. 10.10 Predicted surface location for: indicative search region as at November 2015 marked with *solid line*



17:01 initialisation with no radar data (red) compared with the 18:02 radar initialisation (white)



17:01 initialisation weighted by proximity to 18:02 radar (red) compared with the 18:02 radar based initialisation (white)

Fig. 10.11 Probability density from early initialisation

paths that are preferred after re-weighting using the primary radar data travel South of the radar point and arrive at the Southern edge of the search region. One would expect that, given further computation, samples would be found between these two that would match the 18:02 radar point better and agree better with the search region.

10.7 Cost Index

Another autopilot mode for control of air speed is Cost Index, as discussed in Sect. 6.2.1. To recap, under this mode the autopilot dynamically selects a speed in order to optimise a criterion which trades fuel consumption and travel time according to the Cost Index value entered by the pilot [34]. An additional experiment was performed in which the aircraft was permitted to enter into the Cost Index mode at a random time, remaining in that mode from that time forward. The Cost Index value was randomly sampled between 0 and 100, and lookup tables provided by Boeing were used to determine a speed based on the sampled value of altitude, and an estimate of the aircraft weight, which varies over time as fuel is expended.

The result of this experiment is shown in Fig. 10.12. The white curve shows the pdf calculated through the model using Mach number, as shown in Fig. 10.2, while the red curve shows the pdf calculated using the Cost Index model. The figure shows that the use of Cost Index does not significantly change the result; the distribution is slightly more compact than in the case of the constant Mach number model. Since either mode is possible, the broader result in Fig. 10.2 is preferred for determining the search region.




Fig. 10.12 Set of paths obtained using Cost Index mode for controlling air speed (*red*), compared to result using the Mach number model (*white*) as shown in Fig. 10.2

10.8 Other Variations

The results presented in this chapter are a summary of the analysis we have undertaken on the accident flight data. They are far from an exhaustive set of the variations that have been examined. Many other changes to the filter parameters and measurement model were explored but none made a significant impact on the search region. Example changes to the filter include: modifying the Jeffreys prior; using fixed average manoeuvre times; changing the incremental time step; resampling the Jefferys prior to emulate a change in aircraft behaviour; moving the initialisation point; modifying the initialisation variances; forcing a prescribed angle after 18:39; and quantising heading to integer degrees. Example changes to the measurement model include: vetoing selected measurements; using only C-channel BFO measurements; increasing and decreasing the BFO assumed noise variance; replacing the BFO with a “true” BFO using known location for the validation flights; including or excluding particular BTO measurements and R600 messages; and cross validation where each measurement is removed in turn to investigate output robustness.

In our earliest attempts at producing a pdf a much simpler aircraft model was used. For example, under the assumption of a single turn, the parameter vector is relatively short and more conventional estimation methods can be used. However, such a model cannot be applied to any of the validation flights and runs the risk of constraining the estimated path with additional assumptions. The fact that the very general model here chooses to make few turns provides information that would not be available from a single turn geometric model.

The factors that do make a significant difference to the output pdf are the assumed spread of Mach number and the end of flight model. The assumed Mach number range covers the speeds feasible to achieve the required flight endurance time. The lower end of this speed range results in the Northern part of the pdf and the higher end of the speed range results in the Southern part. Restricting the speed to only Mach numbers above 0.8, for example, would contract the pdf to the South. The consequence is that using a smaller speed range within the bounds already modelled leads to a subset of the search zone. If a different end of flight model is assumed the general consequence is to spread the search zone over a larger area. Simulations have predicted that the maximum distance that the aircraft could have glided under human control is approximately 100nm after 00:19 [5]. The search zone that this scenario would imply is very much larger.

Open Access This chapter is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits any noncommercial use, duplication, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, a link is provided to the Creative Commons license and any changes made are indicated.

The images or other third party material in this chapter are included in the work’s Creative Commons license, unless indicated otherwise in the credit line; if such material is not included in the work’s Creative Commons license and the respective action is not permitted by statutory regulation, users will need to obtain permission from the license holder to duplicate, adapt or reproduce the material.