A Multi-Modal Unscented Kalman Filter for Inference of Aircraft Position and Taxi Mode from Surface Surveillance Data

Harshad Khadilkar* and Hamsa Balakrishnan†

Massachusetts Institute of Technology, Cambridge, MA 02139, USA

We describe a multi-modal unscented Kalman filter developed for estimation of aircraft position, velocity and heading from noisy surface surveillance data. The raw data is composed of tracks generated by the Airport Surface Detection Equipment, Model-X at Boston Logan International Airport, and is obtained from the Runway Status Lights system. The multi-modal filter formulation facilitates estimation of aircraft taxi mode, described by different acceleration and turn rate values, in addition to aircraft states.

Nomenclature

ASDE-X	Airport Surface Detection Equipment, Model-X
MMF	Multi-Modal (Unscented Kalman) Filter
UKF	Unscented Kalman Filter
X	Position (meters east)
Y	Position (meters north)
V	Taxi speed (m/s)
θ	Heading (degrees)
a_f	Acceleration value (m/s^2) in process model of filter f
ω_f	Turn rate value (deg/s) in process model of filter f
$\bar{w}$	Process noise vector
$\bar{v}$	Measurement noise vector
σ^2_{wA}	Assumed process noise variance for any quantity A
σ^2_{vA}	Assumed measurement noise variance for any quantity A
P	Sigma point generator matrix
Q	Process Noise Covariance Matrix
Q_d	Discrete-Equivalent Process Noise Covariance Matrix
R	Measurement Noise Covariance Matrix
$w_{c,i}$	Sigma-point weights for covariance update
$w_{m,i}$	Sigma-point weights for mean update
$\hat{\bar{x}}$	Estimated state vector
$\bar{y}$	Actual Measurement vector
$\hat{\bar{y}}$	Expected Measurement vector
$\bar{x}^{(i)}$	i^{th} sigma point

*PhD candidate, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139. harshadk@mit.edu. AIAA Student Member.

†Assistant Professor, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA 02139. hamsa@mit.edu. AIAA Member.

I. Introduction

A. Motivation

The availability of new surface surveillance data from the Airport Surface Detection Equipment, Model-X (ASDE-X) has created opportunities for detailed investigation of airport surface operations. Specifically, it is now possible to visualize actual taxi routes of aircraft and to directly measure metrics such as total taxi distance, taxi time and time spent in the departure queue. These quantities form the basis for further studies such as the calculation of total fuel burn due to airport surface operations[1,2] or the prediction of aircraft taxi-out times.[3]

However, estimates produced using raw ASDE-X tracks are prone to error, because of inherent noise in the data. For example, noise in position data can lead to overestimation of taxi distance, noise in velocity data can lead to errors in detection of takeoff and landing, and noise in heading data can lead to false turn detection. In order to mitigate these effects, we used an unscented Kalman filter (UKF)[4,5] in the multi-modal form. A bank of filters processes the raw data, with each filter representing a different taxi mode. The current taxi mode is detected by selecting the filter with the least *normalized squared error*, combined with a heuristic that offers resistance to rapid mode jumping. The multi-modal formulation has been previously used in areas such as computer vision,[6] but as far as we know, it has not been used before for aircraft position tracking. In this paper, we present the results from our version of the multi-modal filter, and compare it to raw data and a single-mode filter.

B. Overview of ASDE-X data

Airport Surface Detection Equipment, Model-X is primarily a safety tool designed to mitigate the risk of runway collisions.[7] It incorporates real-time tracking of aircraft on the surface to detect potential conflicts. There is potential, however, to use the data generated by it for surface operations analysis and aircraft behavior prediction. The data itself is generated by sensor fusion from three primary sources: (i) Surface movement radar, (ii) Multilateration using onboard transponders, and (iii) GPS-enabled Automatic Dependent Surveillance Broadcast (ADS-B). Reported parameters include each aircraft's position, velocity, altitude and heading. The update rate is of the order of 1 second for each individual track. The coverage range is approximately 20 miles out from the airport.

ASDE-X data from the Runway Status Lights (RWSL) system at Boston Logan International Airport (BOS) was available for this study. The results presented in this paper are based on data from August to December 2010. As discussed above, it contains a substantial amount of noise despite an upgrade installed in April 2010 to the system at Boston.

II. Formulation

A. Selection of Filter Algorithm

Prior to settling on the UKF as the algorithm of choice, we also tested the linear and extended versions of the Kalman filter.[8] The linear filter was unable to reduce noise to a satisfactory level without sacrificing accuracy, while the extended Kalman filter was prone to divergence. The unscented Kalman filter, on the other hand, was seen to be much more stable. Finally, the multi-modal formulation was used because of its tracking accuracy and the relative ease of inferring taxi mode, when compared to a single filter. While 'truth' was not available in this study since the data was from actual operations, some idea of taxi path could be inferred from the raw trajectories superimposed over a layout of the taxiways. Specifically, aircraft do not stray far from the centerlines marked on the taxiways/runways, allowing comparison of filter tracking behavior by visual inspection.

B. System Dynamics

The state vector to be estimated consists of aircraft position (X in meters east and Y in meters north), velocity (V, m/s) and heading (θ, degrees),

$$\bar{x} = \begin{bmatrix} X & Y & V & \theta \end{bmatrix}^T. \tag{1}$$

The number of states is $n = 4$. To account for variability in longitudinal acceleration and turn rate values, the velocity and heading states were assumed to have independent, zero mean, white Gaussian process noise of variance σ_{wV}^2 and $\sigma_{w\theta}^2$ respectively. The assumption of independence meant that the process noise covariance matrix Q was diagonal,

$$Q = \begin{bmatrix} \sigma_{wV}^2 & 0 \\ 0 & \sigma_{w\theta}^2 \end{bmatrix}. \tag{2}$$

Since all four states are available in ASDE-X data, the measurement vector $\bar{y}$ had the same components as the state vector. It was assumed that the noise in measurement of X and Y positions was independent but had the same variance. The values of position, velocity and turn-rate noise variance were assumed to be equal to σ_{vxy}^2, σ_{vV}^2 and $\sigma_{v\theta}^2$ respectively. The measurement noise covariance matrix, R was therefore given by,

$$R = \begin{bmatrix} \sigma_{vxy}^2 & 0 & 0 & 0 \\ 0 & \sigma_{vxy}^2 & 0 & 0 \\ 0 & 0 & \sigma_{vV}^2 & 0 \\ 0 & 0 & 0 & \sigma_{v\theta}^2 \end{bmatrix}. \tag{3}$$

Finally, the continuous nonlinear state dynamics were derived from kinematic relations to be,

$$\dot{\bar{x}} = \begin{bmatrix} \dot{X} \\ \dot{Y} \\ \dot{V} \\ \dot{\theta} \end{bmatrix} = \begin{bmatrix} V\sin\theta \\ V\cos\theta \\ a \\ \omega \end{bmatrix} + B_w\bar{w}. \tag{4}$$

Here, a is the longitudinal acceleration and ω is the turn rate, both of which are assigned different constant values in the system model for each mode. Since a heading of zero degrees is due north, the sine component of velocity affects X position while the cosine component affects Y position. A full list of taxi modes is given in Table 1. For the single-mode UKF used for comparison purposes, the assumed values were $a = \omega = 0$. The additive noise term in (4) is the product of the the matrix B_w and the process noise $\bar{w}$. The noise vector $\bar{w} = [w_V \quad w_\theta]^T$ contains noise in acceleration and turn rate, and is assumed to vary according to the normal distribution $\mathcal{N}(\mathbf{0}, Q)$. The matrix B_w ensures that $\bar{w}$ affects only velocity and turn rate, and is given by,

$$B_w = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}. \tag{5}$$

The measurement vector is equal to the state vector with an additive noise term,

$$\bar{y} = \bar{x} + \bar{v}. \tag{6}$$

Here, $\bar{v} = [v_X \quad v_Y \quad v_V \quad v_\theta]^T$ is the measurement noise vector, assumed to be independent of process noise, white, Gaussian and to vary like $\mathcal{N}(\mathbf{0}, R)$. The diagonal covariance matrix R ensures that the elements of measurement noise are independent of each other.

C. Filter Implementation

1. *Definition of Constants*

The UKF formulation used for each mode was based on the method proposed by Wan and van der Merwe.[4] According to this method, a state vector of length n is represented by $(2n+1)$ sigma points. Some additional required constants are defined below. Note that in this case, the number of states is $n = 4$.

$$\kappa = 3 - n, \tag{7a}$$
$$\alpha = 0.5, \tag{7b}$$
$$\lambda = \alpha^2(n + \kappa) - n, \tag{7c}$$
$$\beta = 2. \tag{7d}$$

Mode Number	Description	Longitudinal Acceleration a (m/s^2)	Turn Rate ω (deg/s)
1	Default	0	0
2	Moderate Acceleration	1	0
3	Moderate Deceleration	-1	0
4	Constant Speed Turn Right	0	10
5	Constant Speed Turn Left	0	-10
6	Accelerating Turn Right	3	10
7	Accelerating Turn Left	3	-10
8	Decelerating Turn Right	-3	10
9	Decelerating Turn Left	-3	-10
10	High Acceleration	3	0
11	High Deceleration	-3	0

Table 1. List of taxi modes

Based on these constants, the weights used for generation of the $(2n+1)$ sigma points from the sigma point generator matrix are given by,

$$w_{c,i} = \frac{1}{2(n+\lambda)} \qquad \text{for} \quad i = 1, 2, ..., 2n, \tag{8a}$$

$$w_{c,2n+1} = \frac{\lambda}{n+\lambda} + 1 - \alpha^2 + \beta, \tag{8b}$$

$$w_{m,j} = \frac{1}{2(n+\lambda)} \qquad \text{for} \quad j = 1, 2, ..., 2n, \tag{8c}$$

$$w_{m,2n+1} = \frac{\lambda}{n+\lambda}. \tag{8d}$$

Finally, the discrete equivalent process noise covariance matrix, Q_d, is based on the approximation (9), with T_s the time difference between two successive measurements. Note that Q_d is of size 4×4.

$$Q_d \approx B_w Q B_w^T T_s. \tag{9}$$

2. Filter Algorithm

Let us define $z_k^{(i)+}$ to be the value of any quantity z after the measurement update at time step k, for the i^{th} sigma point where $i = 1, 2, ..., 2n+1$. Similarly, $z_{k+1}^{(i)-}$ is the quantity z propagated to time step $k+1$ but prior to the measurement update at time step $k+1$. Let $\hat{\bar{x}}$ denote the estimated state vector, $\bar{x}^{(i)}$ denote the i^{th} sigma point, $\bar{y}$ denote actual measurements, $\hat{\bar{y}}$ be the expected measurement vector and P be the sigma point generator matrix, which is the uncertainty matrix from which sigma points are generated. The filter algorithm[4] is given below. Since all the states are available for measurement, algorithm initialization is done using $\hat{\bar{x}}_0^+ = \bar{y}_0$ and $P_0^+ = R$.

1. Generation of sigma points at time step k, where the i^{th} and $(n+i)^{th}$ sigma points are generated using the transpose of the i^{th} row belonging to the matrix square root of $(n+\lambda)P_k^+$:

$$\begin{aligned}
\bar{x}_k^{(i)+} &= \hat{\bar{x}}_k^+ + [\{\sqrt{(n+\lambda)P_k^+}\}^{(i)}]^T, \quad \text{for} \quad i = 1, 2, ..., n, \\
&= \hat{\bar{x}}_k^+ - [\{\sqrt{(n+\lambda)P_k^+}\}^{(i)}]^T, \quad \text{for} \quad i = (n+1), (n+2), ..., 2n, \\
&= \hat{\bar{x}}_k^+, \quad \text{for} \quad i = 2n+1.
\end{aligned}$$

2. Propagation of belief to the next time step, with a_f and ω_f being the acceleration and turn rate values assumed in the system model for filter f. The propagation equations are based on simple kinematics.

$$\bar{x}_{k+1}^{(i)-} = \bar{x}_k^{(i)+} + \begin{bmatrix} (V_k^{(i)+}T_s + \frac{1}{2}a_f T_s^2)\sin(\frac{\pi}{180}\theta_k^{(i)+}) \\ (V_k^{(i)+}T_s + \frac{1}{2}a_f T_s^2)\cos(\frac{\pi}{180}\theta_k^{(i)+}) \\ a_f T_s \\ \omega_f T_s \end{bmatrix}.$$

3. Calculation of expected state vector at step $(k+1)$ from the propagated sigma points:

$$\hat{\bar{x}}_{k+1}^{-} = \sum_{i=1}^{2n+1} w_{m,i}\bar{x}_{k+1}^{(i)-}.$$

4. Estimation of the uncertainty covariance matrix:

$$P_{k+1}^{-} = Q_d + \sum_{i=1}^{2n+1} w_{c,i}(\bar{x}_{k+1}^{(i)-} - \hat{\bar{x}}_{k+1}^{-})(\bar{x}_{k+1}^{(i)-} - \hat{\bar{x}}_{k+1}^{-})^T.$$

5. Recalculation of sigma points $\bar{x}_{k+1}^{(i)-}$ using the same equations as in step 1, but with P_{k+1}^{-} as the generating matrix. The measurement sigma points, $\bar{y}_{k+1}^{(i)}$, are equal to $\bar{x}_{k+1}^{(i)-}$ since full state measurement is available.

6. Calculation of expected measurement vector at step $k+1$:

$$\hat{\bar{y}}_{k+1} = \sum_{i=1}^{2n+1} w_{m,i}\bar{y}_{k+1}^{(i)}.$$

7. Calculation of the Kalman filter gain at time step $k+1$, L_{k+1}:

$$P_{yy} = R + \sum_{i=1}^{2n+1} w_{c,i}(\bar{y}_{k+1}^{(i)} - \hat{\bar{y}}_{k+1})(\bar{y}_{k+1}^{(i)} - \hat{\bar{y}}_{k+1})^T,$$

$$P_{xy} = \sum_{i=1}^{2n+1} w_{c,i}(\bar{x}_{k+1}^{(i)-} - \hat{\bar{x}}_{k+1}^{-})(\bar{y}_{k+1}^{(i)} - \hat{\bar{y}}_{k+1})^T,$$

$$L_{k+1} = P_{xy}P_{yy}^{-1}.$$

8. Belief update using actual measurement at time step $k+1$, $\bar{y}_{k+1}$:

$$\hat{\bar{x}}_{k+1}^{+} = \hat{\bar{x}}_{k+1}^{-} + L \cdot (\bar{y}_{k+1} - \hat{\bar{y}}_{k+1}),$$

$$P_{k+1}^{+} = P_{k+1}^{-} - LP_{yy}L^T.$$

This algorithm is run simultaneously on all filters in the bank. The expected measurement from each filter is compared to the actual measurement vector. Based on this comparison, the current taxi mode is set, and the state estimate from the corresponding filter becomes the new state estimate for the entire bank. The algorithm then repeats, starting from step 1.

3. Selection of Taxi Mode

The selection of taxi mode is based upon a *Quality Index* (QI) that is essentially the negative of the second-order norm of the residual for each filter, normalized by the measurement noise covariance. Large negative values of QI for a filter f indicate large differences in the actual measurements and the expected ones as derived from the system model for f, and are thus an indicator of poor performance from this filter.

$$\text{QI} = -(\bar{y}_{k+1} - \hat{\bar{y}}_{k+1})^T R^{-1} (\bar{y}_{k+1} - \hat{\bar{y}}_{k+1}). \tag{10}$$

The filter selection algorithm tracks the Quality Index for each filter, for the latest three time steps. The particular filter selected is the one that has the *maximum* (least negative) QI, among all $(2n+1)$ filters, for the *most* times in these three time steps. If three different filters have one instance each of maximum QI in the last three time steps, the incumbent filter continues to be selected. The algorithm is initiated at the start of each flight track with the 'straight, constant velocity' mode active.

A Gaussian likelihood estimate, based on the measurement error of each filter, and the uncertainty as derived from the matrix P, was initially tested as the criterion for mode selection. However, it was found that the measurement noise can cause the error to be occasionally so large that all likelihood values drop to machine zero. In addition, the evaluation of Gaussian likelihood requires a call to a separate subroutine in the code, and is computationally expensive. Therefore, the second-order norm was chosen as a measure of likelihood. The additional heuristic of choosing the best filter over the last three time steps helps avoid frequent mode jumping caused by erratic measurement errors.

III. Results

A. Multi-Modal Filter Output

A comparison of the output from the multi-modal filter with raw data is shown in Figure 1. It is from a portion of taxiway close to the threshold of Runway 33L at Boston Logan. The white line in the upper plot traces the raw data, while the multi-colored segmented line in the lower plot is the filter output. It can be seen that most of the noise in the raw data has been filtered out by the MMF. At the same time, an idea about tracking performance of the filter can be obtained by observing that the filtered output is fairly close to the taxiway markings on the surface. Note that the colors in the filter output indicate the taxi mode as detected by the filter. Table 2 relates the colors to taxi modes. Figure 2 shows sample filtered output for a departure and an arrival. The arriving flight can be seen braking on the runway (in red), prior to turning off onto a rapid-exit taxiway. The departing flight can be seen to be turning onto the runway (portions in light blue) and then accelerating on its takeoff roll (in green).

Figure 3 shows results from the taxi mode detection algorithm for a specific flight. It can be seen that the code correctly detects a constant speed turn (starting at 756 seconds on the X axis) followed by takeoff, which is a period of sustained high acceleration. Note that the fourth plot in the figure shows *relative* Quality Index values for each mode, i.e., the difference between the Quality Index of each filter and the instantaneous Quality Index of the default mode having $[a, \omega] = [0, 0]$. This was done only for ease of plotting.

Figures 4(a) and 4(b) are typical time series plots of the noise in raw data. In Figure 4(a), considerable jitter is seen in the measured Y position, most of which is removed by the filter. There is some noise in the X position as well, but not much in the velocity and heading fields. In Figure 4(b), there is noise in velocity and heading but not in position. Note that each of the four fields is independently measured by the ASDE-X system, and therefore, measured velocity is not simply the derivative of position. Finally, Figure 5 shows autocorrelation plots of the estimation residuals for a single flight. For each measurement variable, it shows the correlation coefficient of the residual (error) at one time step with all other residuals in the rest of the flight track. The low value of correlation except where the difference in index goes to zero (correlation with itself) indicates that the residuals are a white sequence,[9,10] and therefore the filter is working efficiently.

Taxi Mode	Mode Numbers	Color
Raw data	-	White
Straight, constant speed	1, 2, 3	Dark Blue
Decelerating straight / turn	8, 9, 11	Red
Accelerating straight / turn	6, 7, 10	Green
Turn, constant speed	4, 5	Light Blue

Table 2. Color coding for Figures 1 and 2

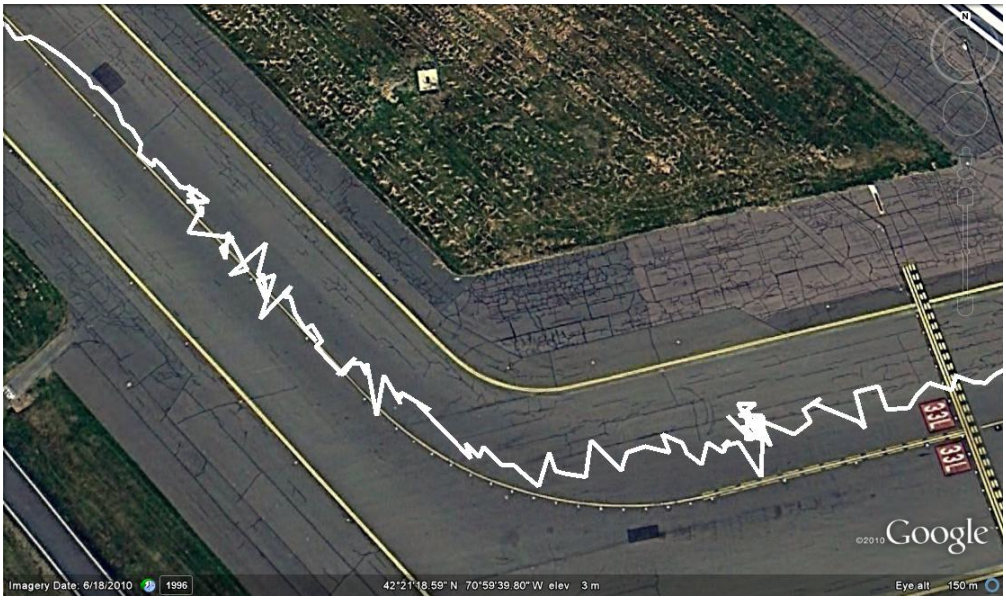

(a) Raw Data

(b) MMF Output

Figure 1. Sample output from Multi-modal Filter. Both pictures correspond to the area enclosed by the orange rectangle in the inset image.

(a) Arriving flight: Landing on the runway is seen on the right, with strong braking marked in red. The flight then leaves the runway and taxis towards its gate beyond the top edge of the picture. Portions where a turn was detected are marked in light blue.

(b) Departing flight: Departure is on Runway 4R, with acceleration phase indicated in green.

Figure 2. Sample MMF output

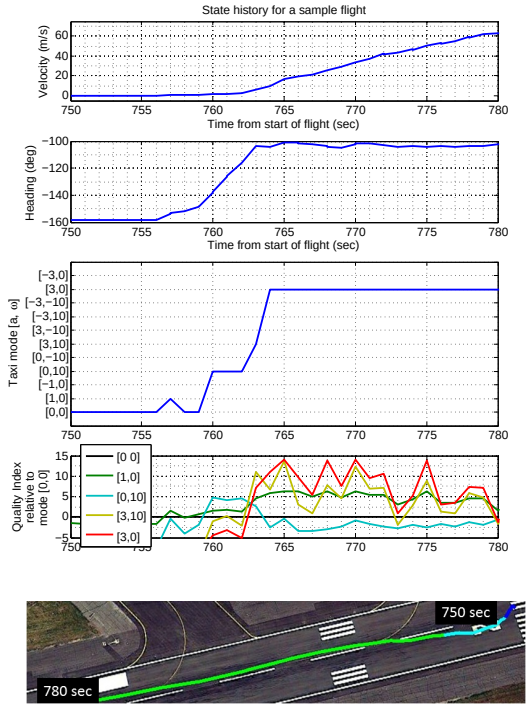

Figure 3. Taxi mode detection for a sample flight. The first two plots show velocity and heading for the part of the aircraft trajectory shown in the bottom plot. The third plot shows the output of the mode detection algorithm, with acceleration and turn-rate values of each mode marked on the Y axis. The fourth plot shows the quality index of the best five modes in this time period, *relative* to the baseline mode $[a, \omega] = [0, 0]$.

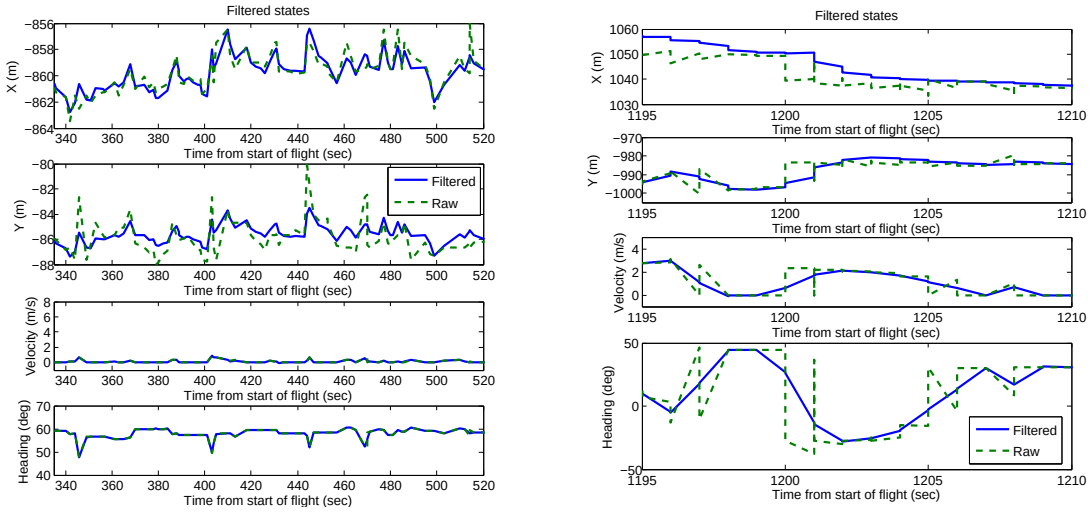

(a) Noise seen in position measurement

(b) Noise seen in heading and velocity measurement

Figure 4. Time series plots of raw and filtered data for sample flights

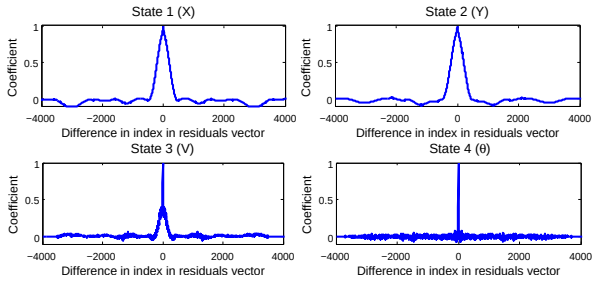

Figure 5. Autocorrelation of residuals for a single sample flight

B. Comparison with single-mode UKF and Raw data

In terms of output quality, the single-mode UKF formulation ends up being inferior to the multi-modal form. This is primarily due to the fact that the process model in the single-mode UKF only accounts for straight, constant speed taxi. If noise suppression comparable to the MMF is desired, the assumed measurement noise covariance needs to be so high that the filter fails to track quick changes in velocity and heading. For example, Figure 6 shows the single-mode filter output belatedly tracking a turning aircraft, thus producing an estimate of position which is unrealistically close to the edge of the taxiway. If certain areas of the airport are being monitored for activity (such as queuing) based on geographical coordinates, this behavior could easily lead to errors.

Figure 7 gives an idea of the relative tracking behavior of the filters discussed so far. It shows the mean ratio of the taxi distance measured using three different methods for a total of 20 flights, to the taxi distance measured visually by a straight-line approximation of the taxi path. The three ways in addition to the visual measurement, by which taxi distances have been calculated are (i) using the MMF, (ii) using the single-mode UKF, and (iii) from raw data. The normalization of taxi distance in this manner accounts for different taxi paths used by different flights. As seen in the figure, the mean distance measured by the MMF is closest to the straight-line method. Note that the MMF output is likely to be closer to actual taxi distance than the straight-line approximation since aircraft do not stick strictly to taxiway centerlines. It is clear from the figure, though, that taxi distance calculated directly using raw data would hold very little value for further analysis, because of its large taxi distance estimates. Further, the large variance precludes the use of constant-factor scaling to increase the accuracy of the raw estimates.

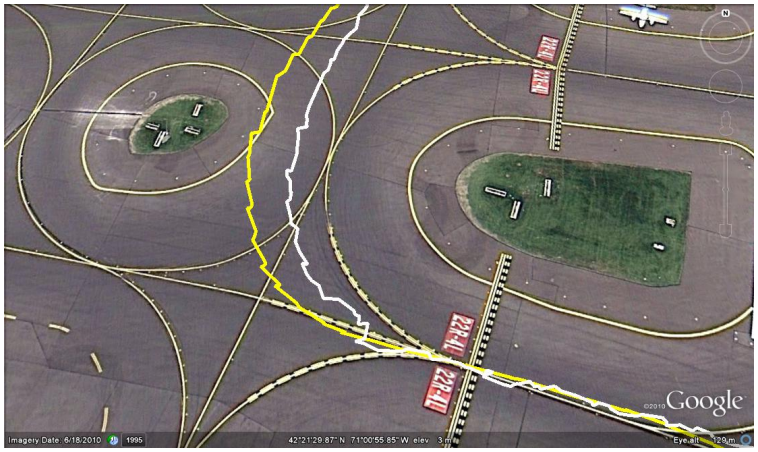

Figure 6. Sample result from single-mode UKF: Raw data in white, UKF output in yellow

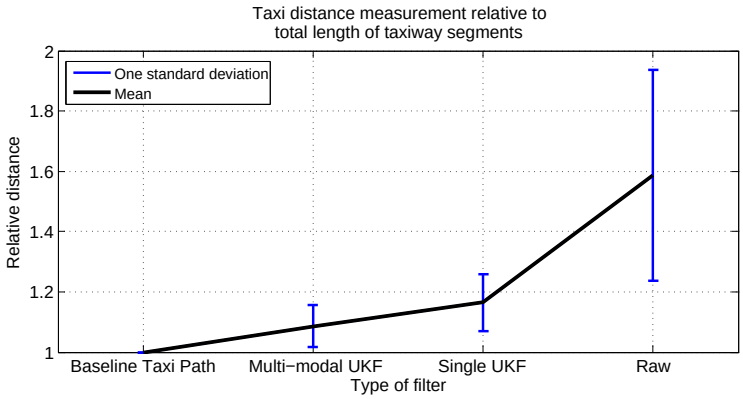

Figure 7. Comparison of taxi-out distance ratios

IV. Conclusion

In this paper, we presented a multi-modal unscented Kalman filter developed to estimate the position, velocity and heading of aircraft on the airport surface. We showed that the raw data contained a significant amount of noise, precluding its direct application to further analysis such as taxi distance calculations. The MMF was shown to successfully mitigate the effect of noise, while maintaining tracking performance. Tradeoffs involved with the use of a single-mode UKF instead of the MMF were considered. It was found that the MMF had better tracking properties than the UKF, making it the filter of choice for applications requiring accurate estimates of aircraft position and taxi mode.

Acknowledgments

We would like to thank Richard Jordan, Mariya Ishutkina and Tom Reynolds from MIT Lincoln Lab for helpful discussions in the context of this work. Thanks also to Jim Eggert and Daniel Herring for help with setting up the ASDE-X data feed.

References

[1]Jung, Y., "Fuel Consumption and Emissions from Airport Taxi Operations," NASA Green Aviation Summit, 2010.

[2]Khadilkar, H. and Balakrishnan, H., "Estimation of Aircraft Taxi-out Fuel Burn using Flight Data Recorder Archives," *AIAA Conference on Guidance, Navigation and Control*, Portland, OR, August 2011.

[3]Jordan, R., Ishutkina, M., and Reynolds, T., "A statistical learning approach to the modeling of aircraft taxi time," *Digital Avionics Systems Conference*, Salt Lake City, UT, October 2010.

[4]Wan, E. and van der Merwe, R., "The Unscented Kalman Filter for Nonlinear Estimation," *Symposium 2000 on Adaptive Systems for Signal Processing, Communication and Control*, IEEE, Alberta, Canada, 2000.

[5]Julier, S. and Uhlmann, J., "A New Extension of the Kalman Filter to Nonlinear Systems," *AeroSense: 11th International Symposium of Aerospace/Defense Sensing, Simulation and Controls*, 1997.

[6]Bellotto, N. and Hu, H., "A bank of unscented Kalman filters for multimodal human perception with mobile service robots," *International Journal of Social Robotics*, Vol. 2, Number 2, 2010.

[7]Sensis Corporation, East Syracuse, NY, *ASDE-X Brochure - ASDE-X 10/05.qxd*, 2008.

[8]Grewal, M. and Andrews, A., *Kalman Filtering: Theory and Practice using MATLAB*, John Wiley & Sons, Inc., 2nd ed., 2001.

[9]Durbin, J. and Watson, G., "Testing for Serial Correlation in Least Squares Regression I," *Biometrika*, Vol. 37, 1950, pp. 409–428.

[10]Durbin, J. and Watson, G., "Testing for Serial Correlation in Least Squares Regression II," *Biometrika*, Vol. 38, 1951, pp. 159–178.