Data-driven Modeling for Aviation Safety Diagnosis and Prognosis

Xiaoge Zhang[1], Yingxiao Kong[2], Abhinav Subramanian[3], and Sankaran Mahadevan[4]

[1,2,3,4] *Department of Civil and Environmental Engineering, Vanderbilt University,
Nashville, TN, 37235, USA*
xiaoge.zhang@vanderbilt.edu
yingxiao.kong@vanderbilt.edu
abhinav.subramanian@vanderbilt.edu
sankaran.mahadevan@vanderbilt.edu

Abstract

The safety of the air transportation system is affected by a variety of uncertainties arising from multiple sources. This paper investigates a diagnosis and prognosis approach to detect anomalies in the flight trajectory, diagnose root causes, and then perform prognosis regarding the risk of occurrence of adverse events, in the presence of various sources of uncertainty. The proposed method is illustrated using a three-step procedure. First, using flight trajectory data, we evaluate the probabilities of system states corresponding to each failure case, from which we formulate a state-space model. Next, we perform anomaly detection for a specific flight trajectory by developing a Bayesian state estimation-based method, and subsequently identify the cause of the detected anomaly. Once the root cause is identified, prognosis is performed to predict the future state in a probabilistic manner. The proposed method is illustrated using near-ground landing data synthetically generated from an open source air traffic simulator – BlueSky. The simulation data mimicking the near-ground landing process with different initial states (e.g., aircraft altitude and speed, response delay, and brake performance) and other factors (such as wind direction) are used to demonstrate the procedures of diagnosis and prognosis.

1. Introduction

As reported by the Federal Aviation Administration (FAA), in 2016, 2.6 million passengers flew every day in and out of U.S. airports, and 39.9 billion pounds of freight were shipped by air (Federal Aviation Administration, 2017). The increasing air travel demand, which is expected to double in two decades, will make the system overloaded and congested; therefore, the safety of the air transportation system has received considerable attention in recent years (Sankararaman, Roychoudhury, Zhang, & Goebel, 2017; X. Zhang & Mahadevan, 2017, 2018), along with other issues related to traffic management (Guimera, Mossa, Turtschi, & Amaral, 2005; Di Gravio, Mancini, Patriarca, & Costantino, 2015).

In the current en route air traffic control system, the aircraft is required to fly along a predesigned route (or trajectory) comprised of a sequence of fixed waypoints that connect the origin and destination airports. One commonly encountered situation in assessing the safety of the air transportation system is the flight trajectory deviating from its filed flight plan. Under this circumstance, it is important to find out what causes the flight trajectory to deviate from its filed flight plan, when the flight trajectory starts to deviate, how the anomalous flight trajectory is going to evolve over the subsequent waypoints, and whether the anomalous flight trajectory will progress to a state that endangers the safety of the current flight and other nearby flights. To answer these questions, it is crucial to develop a real-time fault detection and identification (FDI) algorithm to perform anomaly detection, fault diagnosis, and system state prognosis. Over the past few years, numerous methods have been developed for fault detection and isolation based on state estimation theory (Zhao, Skjetne, Blanke, & Dukan, 2014; Kadirkamanathan, Li, Jaward, & Fabri, 2002), generalized likelihood ratio tests (Kadirkamanathan et al., 2002), model-based approach (Yin & Zhu, 2015), among others (B. Zhang et al., 2011). For example, Li & Kadirkamanathan (Li & Kadirkamanathan, 2001) combined likelihood ratio with particle filter for fault detection and isolation in stochastic nonlinear systems; Yin and Zhu (Yin & Zhu, 2015) integrated the features in genetic algorithm with particle filter to implement an intelligent algorithm for detecting the system faults in real time for a nonlinear system.

In the context of the air transportation system, few studies have been conducted related to near-ground flight trajectory anomaly detection, diagnosis, and prognosis. A recent study by Di Ciccio et al. (Di Ciccio, Van der Aa, Cabanillas, Mendling, & Prescher, 2016) trained a one-class support vector machine-based classification model to detect flight trajectory anomalies based on several features extracted from flight track events data. However, they only address the anomaly detection problem while several other important issues remain to be answered, such as why the flight trajectory deviates from its flight plan, and how the anomalous flight trajectory will evolve given the diagnosed system malfunction or other causes. In this paper, we are motivated to fill this gap by developing a diagnosis and prognosis framework. To accomplish this goal, we utilize an open source air traffic simulator – BlueSky (Hoekstra & Ellerbroek, 2016)– to mimic the near-ground landing process, in which different types of anomalies and their combined effects are simulated. The abnormal behavior during landing could be caused by many different reasons, such as initial approach characteristics, communication and response delay, brake failure, tailwind, crosswind, etc.

For the sake of illustration, we simulate three different system fault modes, namely response delay, brake failure, and high-speed approach, which result in the different failure scenarios of aircraft landing. By using the simulation data collected from BlueSky, we develop a systematic data-driven diagnosis and prognosis framework with the following features:

- Fault detection and fault diagnosis: A Bayesian state estimation-based approach is developed to detect faults, where the system state is modeled as the aircraft position represented by latitude and longitude at a specific altitude. With flight landing data, a probabilistic approach is developed to estimate the conditional probability of anomaly occurrence at a given altitude under each possible fault mode. When the anomaly in terms of aircraft position is detected (based on a specified threshold), then the fault mode is determined based on the conditional probability information.
- Flight trajectory prognosis: Given the diagnosed root cause, prognosis is performed to make prediction of aircraft position along the subsequent altitudes. The early estimation of flight trajectory evolution over time enables the assessment of system safety in advance before it progresses to a state that impacts system safety.

The remainder of the paper is organized as follows. In Section 2, we introduce the necessary background information and formulate the problem of interest. In Section 3, we introduce the methods that are used to perform diagnosis and prognosis. In Section 4, we present the numerical results and demonstrate the performance of the developed method. In Section 5, we provide concluding remarks.

2. PROBLEM STATEMENT

The problem of anomaly detection consists of making decision on the presence or absence of anomaly in the monitored system, while fault diagnosis refers to identifying the active fault mode among a number of possible modes. In this paper, we assume that the data supplied by the flight system is the periodic update on the location of the aircraft during landing. When the flight reaches each target altitude, its position information is reported to the ground traffic controller, from which we determine whether anomalous behavior is present, and if so, what causes the anomaly, and how the diagnosed cause will influence the flight position over the subsequent target altitudes, and whether the aircraft will land safely on the runway. Throughout the paper, we assume that normal behavior and all possible faulty behaviors of the flight system can be described by a given finite set of state-space models indexed by $m = 0, 1 \cdots, M$:

$$x_k = f_{k-1}^{(m)}\left(x_{k-1}, w\right); k = 1, 2, \cdots \quad (1)$$

$$z_k = h_k\left(x_k, v_k\right); k = 1, 2, \cdots \quad (2)$$

where k is an index representing the aircraft's altitude, x is the 2×1 state vector consisting of the latitude and longitude of the aircraft position, w is the vector of random parameters (namely the initial speed, response delay, and wind direction), $f^{(m)}$ is the state transition function associated with fault mode $m = 0, 1 \cdots M$, z represents the vector of measurement data, and v is the measurement noise vector.

Definition 2.1. Anomaly detection refers to the detection of a shift from the normal mode ($m = 0$) to a faulty mode ($m = 1 \cdots, M$).

Definition 2.2. Fault diagnosis is deciding which of the M faulty models the system has shifted to.

In this paper, we receive periodic measurement data on the location (latitude and longitude) of the flight at each altitude from the air traffic simulator BlueSky. A data-driven approach will be developed in the next section for anomaly detection and fault identification.

3. PROPOSED METHOD

In this section, we demonstrate the proposed method for flight behavior anomaly detection, fault diagnosis, and flight trajectory prognosis.

3.1. Anomaly Detection

We present a Bayesian state estimation-based approach to anomaly detection, using Eqs. 1 and 2 respectively as the process and measurement models.

Let the anomaly be represented by the variable χ, which can assume values $0, 1, \cdots M$, with 0 representing the normal mode, and $1, 2, \cdots M$ representing the M fault modes. $P\left[\chi_k = m\right]$, $m = 1, 2, \cdots M$ represents the probability of anomalous behaviour of mode m at altitude index k, and $P\left[\chi_k = 0\right]$ represents the probability of the absence of anomaly. These probabilities, conditioned upon the measured data, can be represented as

$$P\left[\chi_k = m|y_{1:k}\right] = \int P\left[\chi_k = m|x_k, y_{1:k}\right] p\left(x_k|y_{1:k}\right) dx_k \qquad m = 1, 2, \cdots \tag{3}$$

In the above equation, $p\left(x_k|y_{1:k}\right)$ can be evaluated using Bayesian state estimation methods.

In the problem of anomaly detection, we are concerned with the first occurrence of an anomaly, whose probability would be given by $P\left[\chi_1 = 0, \chi_2 = 0, \cdots, \chi_{k-1} = 0, \chi_k = m\right]$, for $m = 1, 2, \cdots, M$. If this probability exceeds a pre-specified tolerance limit for a certain fault mode m at an altitude index k, we can conclude that an anomaly of type m has been detected.

3.2. Fault Diagnosis

In the proposed methodology, fault diagnosis is performed alongside anomaly detection. Using the relations presented in the preceding section, we obtain the probability of occurrence of anomaly of type m at the altitude index k as $P\left[\chi_k = m|y_{1:k}\right]$. In this study, a particle filtering approach is adopted to evaluate these probabilities. We begin with a data set composed of N_m samples of flight trajectories corresponding to the $m = 1, 2, \cdots, M$ fault modes and the one normal mode. These are treated as samples for the standard particle filtering algorithm, and $P\left[\chi_k = m|y_{1:k}\right]$ is evaluated for each m as

$$\begin{aligned} P\left[\chi_k = m|y_{1:k}\right] &= \int P\left[\chi_k = m|x_k, y_{1:k}\right] p\left(x_k|y_{1:k}\right) dx_k \\ &\propto \sum_{i=1}^{N_m} \prod_{r=1}^{k} p\left(y_r|x_{m,r}^{(i)}\right) \end{aligned} \tag{4}$$

Here $x_{m,r}^{(i)}$ represents the flight location corresponding to the $i-th$ sample, at altitude index r, for the fault mode m. Following the detection of an anomaly, the probabilities $P\left[\chi_k = m|y_{1:k}\right]$ are evaluated for $m = 1, 2, \cdots, M$, and if the source of the anomaly is identifiable using the given data, the probability corresponding to one of the modes $m = 1, 2, \cdots, M$ would be comparatively higher. This is analogous to the problem of model selection.

3.3. Flight Trajectory Prognosis

After the fault mode is identified, we seek to predict the effect of the system fault on the future trajectory. If a flight is diagnosed as having fault of mode m at altitude index k, its position at subsequent altitudes can be estimated using the equation:

$$x_r = x_k + \Delta x_{r,k}^m; r > k \tag{5}$$

where $\Delta x_{r,k}^m$ is estimated by resampling from the distribution corresponding to the data set for the $m-th$ fault mode based on the updated state x_k.

4. NUMERICAL EXAMPLE

We illustrate the proposed methodology for a near-ground landing process in an airport. The runway's angle of the direction from north is 135.2 degrees. The position of the aircraft at four respective altitudes (700 meters, 500 meters, 300 meters and 100 meters) above the airport is measured, and analyzed for anomaly detection, fault diagnosis, and landing prognosis.

The flight trajectory data is generated using BlueSky simulations. In BlueSky, the wind can be modeled by defining a wind vector (wind direction and magnitude) at a specific position and altitude. Different wind vectors can be defined at different locations. To simulate the impact of wind direction on aircraft landing, we consider eight different wind directions, and their angles of the direction from the north are: -44.80, 0.20, 45.20, 90.20, 135.20, 180.20, 225.20, and 270.20 (unit: degree). In particular, the wind with the direction of -44.80 degree is referred to as tail wind, while the wind with the direction of 135.20 represents the head wind. The second factor considered in this example is the delay in communication and pilot response, which can significantly affect the safety of landing. A quantitative measure of delay in reducing the aircraft speed is used in BlueSky simulation to represent this effect with three values: 10 seconds, 30 seconds and 50 seconds, and the normal case is when there is no delay. Thirdly, the aircraft approach speed is also taken into account, with two values for illustration: normal approach speed (150 knots), and anomalous approach speed (180 knots). The fourth factor relates to the brake performance in reducing the aircraft speed once it touches down, with three deceleration rates: 0.8 m/s^2, 1.131 m/s^2, and 2.2 m/s^2 (normal). The anomalous deceleration values may be caused either by brake system failure or by slickness of the runway under rainy or snowy conditions.

Table 1 summarizes the aforementioned variables and their values during aircraft landing. The last column of Table 1 reports the variable settings in the normal case. For example, in the normal case, aircraft approaches the near terminal area at a speed of 150 knots without any response delay. After it touches down, the brake should be able to decelerate the aircraft at a rate of 2.2 m/s^2. The third column shows the considered values, indicating individual fault mode pertinent to each variable. When only single isolated fault modes are considered, only the variable related to that fault mode has abnormal value while all the remaining variables take normal values. To account for the variability of each individual

Table 1. The basic configuration of BlueSky simulation

Variable Name	Distribution	Mean	Standard Deviation	Normal Value
Speed (unit: knot)	Normal	150, 180	10	150
Response delay (unit: second)	Normal	10, 30, 50	2	0
Wind direction (unit: degree)	Normal	-44.8, 0.20, 45.2 90.2, 135.2, 180.2 225.2, 270.2	2	–
Brake (unit: m/s^2)	–	-0.8, -1.131, -2.2	–	-1.131

Table 2. Failure probability of each considered case

Fault Mode	**1**	**2**	**3**	**4**	**5**	**6**
	speed: 150 knots	**speed**: 180 knots	**response delay**: 10 seconds	**response delay**: 30 seconds	**response delay**: 50 seconds	**wind direction**: -44.8
Fault Mode	**7**	**8**	**9**	**10**	**11**	**12**
	wind direction: 0.2	**wind direction**: 45.2	**wind direction**: 90.2	**wind direction**: 135.2	**wind direction**: 180.2	**wind direction**: 225.2
Fault Mode	**13**	**14**	**15**	**16**		
	wind direction: 270.2	**combined fault 1**: wind direction: 45.2 response delay: 10 secs	**combined fault 2**: wind direction: 45.2 response delay: 30 secs	**combined fault 3**: wind direction: 45.2 response delay: 50 secs		

fault mode, we run 100 simulations for each considered fault mode. In addition to single isolated fault modes, simulations are also performed to analyze the combined effects of several hybrid fault modes. Table 2 lists the fault configuration for all the considered cases. The first configuration in Table 2 represents normal operation with no faults, with a speed of 150 knots, no wind, no delay in response, and no brake failure (i.e., ground deceleration of 2.2 m/s^2)

Fig. 1 displays the latitude distribution at different altitudes across all the simulation cases, where the blue segment denotes the normal case, while yellow segments represent all the other anomalous cases. It can be observed that each fault mode results in a different latitude distribution. Comparison of this distribution against the normal case is useful in the subsequent diagnosis of the root cause by measuring the degree of aircraft latitude deviation away from normal case.

Fig. 2 shows the touchdown locations of the aircraft on the ground across all the simulation cases, and each color corresponds to one fault mode. In Fig. 2, the black solid line denotes the runway that is selected for the simulation, with the beginning and ending of the runway shown by two arrows. As can be observed, in some configurations, the aircraft fails to land safely on the runway. It is observed some of the flights land ahead of the runway as shown in the right bottom of Fig. 2, while some flights have incorrect lineup with the runway after landing, i.e., they land on either side of the runway.

Fig. 3 shows the distribution of final aircraft position among the aircraft that correctly land on the runway under different deceleration scenarios. Note that compared to Fig. 2 which shows the touch down position of all aircraft, Fig. 3 only contains the flight that touches down on the runway correctly, i.e., the flights with incorrect lineup are ignored. This implies that touch down outside the runway is considered a failure, and among those flights that land on the runway, Fig. 3 is analyzing for additional failure, namely runway over-run due to failure of the braking system. Note that even there is no brake failure, some aircraft might experience runway overrun due to high touchdown speed, as observed in Fig. 3(c).

Following the method introduced in Section 3, we have tested the performance of the proposed approach in correctly detecting the system anomaly and diagnosing the root cause by splitting the dataset into two parts: model generation and performance test, and the same process is repeated for three times as cross-validation. Table 3 presents the statistical results of the algorithm in anomaly detection and fault diagnosis given different amounts of observed data, i.e., the data at different altitudes. At each altitude, we receive a new observed flight position data, anomaly detection and diagnosis is performed again. In Table 3, the first column denotes algorithm accuracy when only the flight position data at the altitude of 700 meters is available, the second column represents the accuracy of the developed approach in anomaly detection and fault identification given the flight position data at the altitude of 700 and 500 meters, and so on. As expected, the algorithm's performance improves with the increase in the observed data. Table 3 reports the accuracy values over the three cross-validation tests.

The above analysis demonstrates the overall performance of the developed algorithm in anomaly detection and fault identification. Next, we consider a specific case (fault mode 13) to illustrate the developed method for flight position prognosis. The Bayesian state estimation method correctly diagnoses the

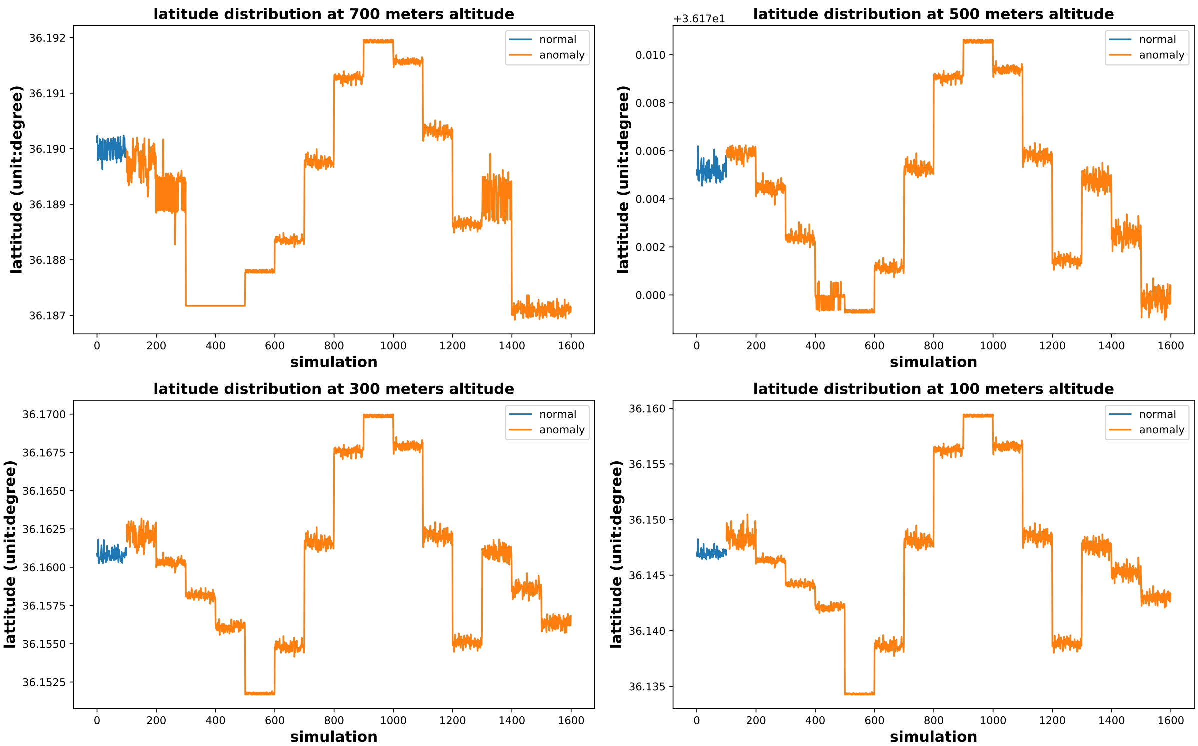

Figure 1. The latitude of all the simulation cases at multiple altitudes

Table 3. Performance evaluation of the proposed method in anomaly detection and fault diagnosis with different amounts of observed data

Algorithm accuracy with observation at 700 meters	Algorithm accuracy with observation at 700 and 500 meters	Algorithm accuracy with observation at 700, 500 and 300 meters	Algorithm accuracy with observation at 700, 500, 300 and 100 meters
0.7033	0.9133	0.9400	0.9567

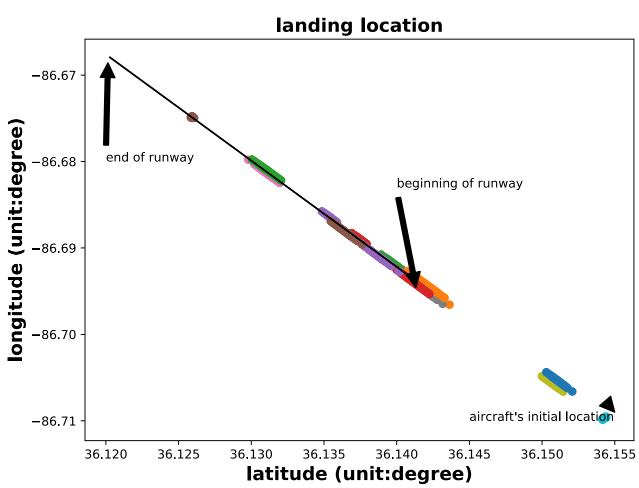

Figure 2. Touchdown locations across all the simulation cases

fault mode with the observed flight position data at the altitude of 700 meters. Based on the diagnosis result, we estimate the flight position along the subsequent altitudes. Fig. 4 shows the prognosis of flight position in terms of latitude and longitude over the subsequent altitudes. In Fig. 4, the blue curve denotes the probabilistic estimation of the flight position using Eq. (5). It can be observed that the actual aircraft latitude and longitude stay within the prior predicted distribution of flight latitude and longitude considering all scenarios. Along with the flight position prediction at subsequent altitudes, we also estimate the probability of failure (i.e., landing outside the runway) for this particular case as 0.736, based on the diagnosis result at 700 meters altitude.

5. CONCLUSION

In this paper, a probabilistic computational algorithm is developed to perform anomaly detection, fault diagnosis and prognosis of flight trajectory. The problem of anomaly detection is treated as one of state estimation, where the system state is the flight position at different altitudes during landing,

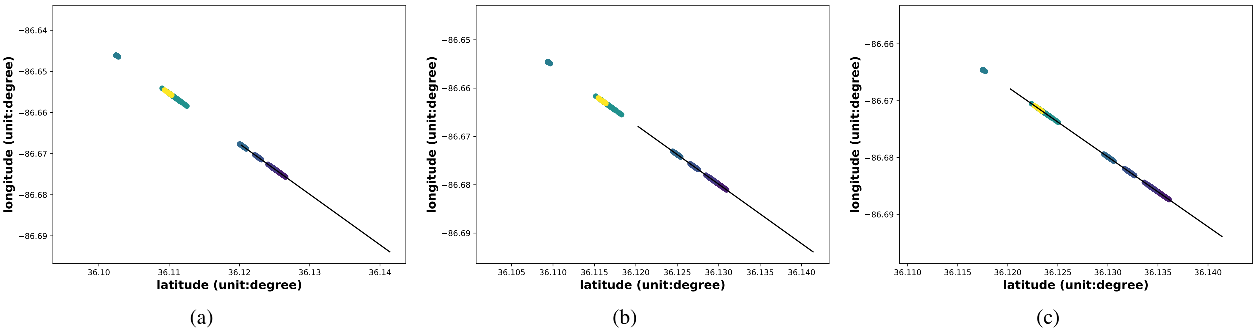

Figure 3. Final aircraft location on the runway after deceleration. (a) -0.8 m/s^2. (b) -1.131 m/s^2. (c) -2.2 m/s^2 (no brake failure)

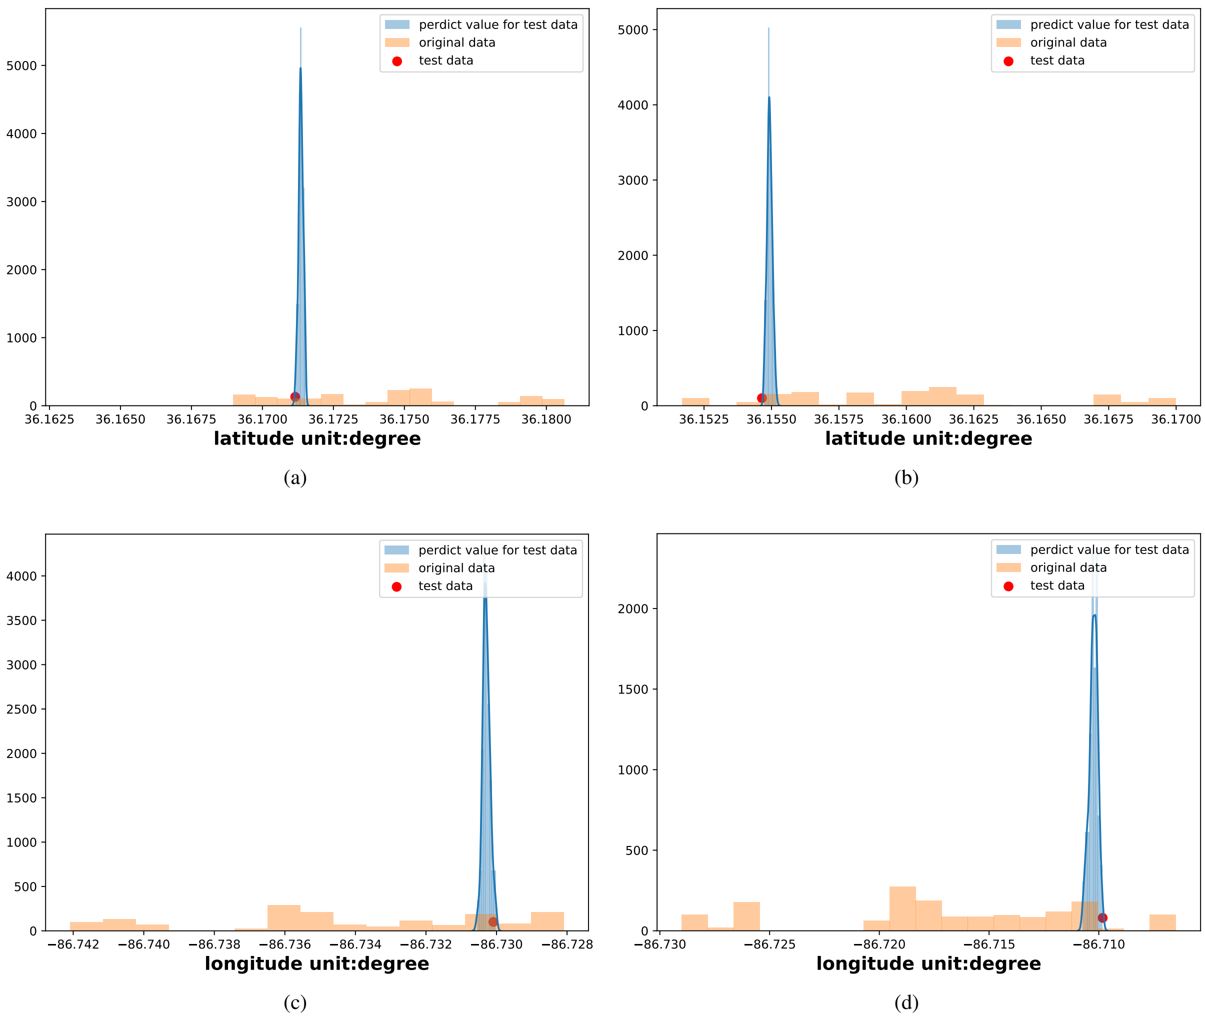

Figure 4. Flight position prognosis based on diagnosis at 700 meters altitude: (a) Latitude at 500 meters. (b) Latitude at 300 meters. (c) Longitude at 500 meters. (d) Longitude at 300 meters

and the various failure cases diagnosed based on observations of these states. This approach permits real-time fault identification, and is able to account for the presence of multiple faults. The methodology can be further extended to systemlevel safety assessment by modeling the interactions among flight trajectories of multiple aircraft if the trajectory information is available.

With an open source air traffic simulator (BlueSky), simulations are performed to mimic the landing process for all the considered sixteen cases (13 single isolated fault modes and three combined fault modes). Additional simulations are also performed to verify the performance of the diagnosis methodology. As illustrated in Section 4, the developed algorithm demonstrates reasonable performance in detecting the anomalous landing trajectory and diagnosing the fault mode corresponding to the observed anomaly, especially as the aircraft position is monitored over several steps.

Although the fault modes considered in this paper are simplified representations of reality, the proposed state-space estimation method is generalizable to account for more realistic scenarios. Also, fault modes such as high approach velocity, response delay and braking system failure represent the effects of other underlying causes, and future work can consider extending this approach to diagnosis through the hierarchy of fault modes.

ACKNOWLEDGMENT

The research reported in this paper was supported by funds from NASA University Leadership Initiative grant (Grant No. NNX17AJ86A, Project Technical Monitor: Dr. Kai Goebel) through subcontract to Arizona State University (Principal Investigator: Dr. Yongming Liu). The support is gratefully acknowledged.

REFERENCES

Di Ciccio, C., Van der Aa, H., Cabanillas, C., Mendling, J., & Prescher, J. (2016). Detecting flight trajectory anomalies and predicting diversions in freight transportation. *Decision Support Systems*, *88*, 1–17.

Di Gravio, G., Mancini, M., Patriarca, R., & Costantino, F. (2015). Overall safety performance of air traffic management system: Forecasting and monitoring. *Safety Science*, *72*, 351–362.

Federal Aviation Administration. (2017). *Air traffic by the numbers,* https://www.faa.gov/air_traffic/by_the_numbers/.

Guimera, R., Mossa, S., Turtschi, A., & Amaral, L. N. (2005). The worldwide air transportation network: Anomalous centrality, community structure, and cities' global roles. *Proceedings of the National Academy of Sciences*, *102*(22), 7794–7799.

Hoekstra, J. M., & Ellerbroek, J. (2016). Bluesky atc simulator project: an open data and open source approach. In *Proceedings of the 7th international conference on research in air transportation* (pp. 1–8).

Kadirkamanathan, V., Li, P., Jaward, M. H., & Fabri, S. G. (2002). Particle filtering-based fault detection in nonlinear stochastic systems. *International Journal of Systems Science*, *33*(4), 259–265.

Li, P., & Kadirkamanathan, V. (2001). Particle filtering based likelihood ratio approach to fault diagnosis in nonlinear stochastic systems. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, *31*(3), 337–343.

Sankararaman, S., Roychoudhury, I., Zhang, X., & Goebel, K. (2017). Preliminary investigation of impact of technological impairment on trajectory-based operations. In *17th AIAA Aviation Technology, Integration, and Operations Conference, AIAA Aviation Forum.*

Yin, S., & Zhu, X. (2015). Intelligent particle filter and its application to fault detection of nonlinear system. *IEEE Transactions on Industrial Electronics*, *62*(6), 3852–3861.

Zhang, B., Sconyers, C., Byington, C., Patrick, R., Orchard, M. E., & Vachtsevanos, G. (2011). A probabilistic fault detection approach: Application to bearing fault detection. *IEEE Transactions on Industrial Electronics*, *58*(5).

Zhang, X., & Mahadevan, S. (2017). Aircraft re-routing optimization and performance assessment under uncertainty. *Decision Support Systems*, *96*, 67–82.

Zhang, X., & Mahadevan, S. (2018). A hybrid data-driven approach to analyze aviation incident reports. In *2018 aviation technology, integration, and operations conference, AIAA AVIATION Forum, (AIAA 2018-3982).*

Zhao, B., Skjetne, R., Blanke, M., & Dukan, F. (2014). Particle filter for fault diagnosis and robust navigation of underwater robot. *IEEE Transactions on Control Systems Technology*, *22*(6), 2399–2407.

BIOGRAPHIES

Xiaoge Zhang received his B.S. degree from Chongqing University of Posts and Telecommunications in 2011, and the M.S. degree from Southwest University in 2014, both in Chongqing, China. Currently, he is pursuing his Ph.D. degree in Vanderbilt University, and is expected to graduate in early 2019. From August to December in 2016, he interned at the National Aeronautics and Space Administration (NASA) Ames Research Center (ARC), Moffett Field, CA, working at the Prognostics Center of Excellence (PCoE) led by Dr. Kai Goebel. He was a recipient of the Chinese Government Award for Outstanding Self-financed Students Abroad in 2017. He has published more than 30 research papers in peer-reviewed journals, such as IEEE Transactions on Cybernetics, IEEE Transactions on Reliability, Decision Support Systems, Information Sciences, Safety Science, Annals of Operations Research, Reliability Engineering and System Safety, and International Journal of Production Research, among others. Four of his publications are Essential Science Indicators (ESI) highly cited papers. His current research interests include uncertainty quantification, reliability assessment, risk analysis, network optimization, and data analytics. He is a student member of IEEE, INFORMS, and SIAM.

Yingxiao Kong received the B.S degree from Tianjin University in 2017, Tianjin, China. Currently, she is pursing her Ph.D degree at Vanderbilt University, Nashville, TN, USA. She has coauthored several international journal and conference papers. Her research interests include simulation-based modeling, system safety assessment, and data analytics.

Abhinav Subramanian is a postdoctoral research scholar at Vanderbilt University, Nashville, TN, USA. He received his BE in Civil Engineering from MVSR Engineering College, Hyderabad, India, and his Ph.D. in Civil Engineering from the Indian Institute of Science, Bangalore, India. His research interests include system identification, structural dynamics, and model uncertainty quantification. He has published several journal and conference papers on these topics.

Sankaran Mahadevan is John R. Murray Sr. Professor of Engineering, and Professor of Civil and Environmental Engineering at Vanderbilt University, Nashville, Tennessee, where he has served since 1988. His research interests are in the areas of uncertainty quantification, model verification and validation, reliability and risk analysis, design optimization, and system health monitoring, with applications to civil, mechanical and aerospace systems. His research has been extensively funded by NSF, NASA, FAA, DOE, DOD, DOT, NIST, GE, GM, Chrysler, Union Pacific, American Railroad Association, and Sandia, Idaho, Los Alamos and Oak Ridge National Laboratories. His research contributions are documented in more than 600 publications, including two textbooks on reliability methods and 280 journal papers. He has directed 42 Ph.D. dissertations and 24 M. S. theses, and has taught several industry short courses on reliability and risk analysis methods. He is a Fellow of AIAA and EMI (ASCE). His awards include the NASA Next Generation Design Tools award (NASA), the SAE Distinguished Probabilistic Methods Educator Award, SEC Faculty Award, and best paper awards in the MORS Journal and the SDM and IMAC conferences. Professor Mahadevan obtained his B.S. from Indian Institute of Technology, Kanpur, M.S. from Rensselaer Polytechnic Institute, Troy, NY, and Ph.D. from Georgia Institute of Technology, Atlanta, GA.